

Cobb's Guide to PC and LAN Security



Stephen Cobb, CISSP

**Electronic Edition
Part One of Three
Chapters 1-5**

**THIS FREE ELECTRONIC EDITION IS LICENSED BY THE AUTHOR FOR USE
UNDER CREATIVE COMMONS, ATTRIBUTION, NON-COMMERCIAL, NO DERIVATES.
Please see <http://creativecommons.org/licenses/by-nc-nd/3.0/>**

This text was converted to text from high resolution page scans using Adobe OCR software. You may encounter OCR errors and artifacts. No claims are made as to accuracy of the information in this document. Use at your own risk and discretion.

Contents

Acknowledgments	xi	
Introduction	xiii	
Why Choose This Book?		xiii
Is This Book for You?		xv
How This Book Evolved		xvi
Where to Find What		xvii
Special Foreword by the Author	xix	
Something Old, Something New		xix
The Human Factor		xx
A Word of Thanks		xx ii
Chapter 1. Security Matters		1
The State We Are In		1
Now the Good News		7
Defining Security		12
Computer Types		12
LANs, WANs , and Other Permutations		16
Categories of Personal Computers		21
User Categories		25
What Is at Stake		26
Attacks, Threats, and Scares		28
Do You Really Need This Book?		29
Further Questions of Security		34
The Network Connection		40
Summary		40
Chapter 2. Security Solutions		42
People, People, People		43
Backup, Backup, Backup		45
Turning on Your Computer		46
Autoexecute Override		57
Secure Boots		58
Under Lock and Key		64
Basic File Protection		70

While You're Away from Me	73
Further Measures	75
The Network Connection	76
Summary	78
Chapter 3. Security Planning	80
Questions, Concepts, and Cycles	80
Basic Risk Analysis	86
Tools and Techniques	93
Security Policy	103
Getting to Grips with Policy	111
Disaster-Recovery Planning	112
Stay Tuned	116
The Security Audit	116
The Network Connection	127
Summary	129
Chapter 4. Secure Hardware	131
Physical Security	131
A Secure Example	136
Exercising Restraint	137
Portables, Alarms, and Other Ideas	144
Tamper Resistance and Identification	148
Securing the Right Thing	154
Key Management	155
Computer Insurance	158
Removable Media	161
The Network Connection	165
Summary	165
Chapter 5. Secure Power	167
Power to the Computer	167
Fuses, Grounds, and Breakers	173
Regulating the Power Supply	180
Sags and Line Conditioning	188
Noise and Static	189
Guaranteeing the Power Supply	193
Software Assistance	209
Electronic Eavesdropping and Worse	210
The Network Connection	213
Summary	214
Chapter 6. Secure Sites	216
Securing the Perimeter	216
Card Control Systems	221
Access and Social Engineering Attacks	230

Smartcards	230
Biometrics	233
System Access Control	242
DOS Access Controls	245
Network Connection	250
Summary	251
 Chapter 7. Secure Access	 254
The Elements of Limited Access	254
The Elements of Control	259
Access Scenarios	266
Introduction to Cryptography	271
DES, RSA, and Today's Encryption	284
Public-Key Encryption	293
Practical Encryption	298
The Password Problem	305
The Network Connection	310
Summary	310
 Chapter 8. Secure Data	 312
Backup Basics	312
Backup Strategies	320
Backing Up on Floppies	333
Compression, Copies, and Archives	336
Backup Hardware: Tape	347
Backup Hardware: Removable Disks	353
Backup Software	361
Software Safety Nets	365
Disk Disaster Recovery and Prevention	367
The Network Connection	377
Summary	378
 Chapter 9. Secure Code	 381
MC: An Instant Intro	381
Plan of Attack	383
Defining the Problem	385
Basic Defenses against MC	396
What Viruses Do	399
Where Does MC Come from?	413
Virus Symptoms	420
Antivirus Tools	420
Virus Response and Recovery	425
Virus Myths and Misconceptions	426
The Network Connection	431
Summary	434

Chapter 10. Secure Software	436
The Big Picture	436
Software Piracy	440
Why Cheat?	448
A Brief History of Copy-Protection	450
Other Rights and Wrongs	456
The Network Connection	463
Summary	463
Chapter 11. Secure Networks I	465
LAN Security: Basics	465
LAN Security: Structure and Goals	470
LAN Functions, Concerns, and Policy	474
LAN Threats	476
LAN Security Services and Mechanisms	482
LAN Hardware Security	490
LAN Fault Tolerance	494
Summary	499
Chapter 12. Secure Networks II	501
Method in the Madness	501
NetWare: Then and Now	502
NetWare 3.x: Security Features	503
NetWare 3.x Tips, Tricks, and Attacks	512
NetWare 4.x	525
NetWare Extras	533
Security and Networking with Windows	538
Securing NTAS Networks	545
Windows Network Warnings	551
Summary	554
Chapter 13. Secure Communications	556
Basic Communications	556
Security and Outbound Calls	563
Problems on Two-way Street	565
Remote-Access Security	568
The Internet Problem	572
Other Forms of Communication	586
Summary	588
Chapter 14. Secure People	590
The People Problem	590
A People Solution	594
People: A Second Opinion	595
Hackers	601
Hacking and the Law	610
Protection against Hacking	611

Social-Engineering Attacks	613
The Ethics Thing	615
Summary	619
Chapter 15. Security in the Future	621
The Layered Approach	621
The System-Centric Approach	627
Positive Signs	634
Mixed Prospects	639
Summary	643
Appendix A. Threat List	645
Basic Categories	645
Functional Categories (after Stallings)	646
Passive Attacks	646
Active Attacks	646
Appendix B. A Brief Guide to Batch Files	647
How Batch Files Work	647
Using Batch Files	650
Appendix C. Computer Security Policy	652
Introduction to Computer Security Policy	652
LAN Security Policy	653
General LAN Policies	655
Specific Responsibilities for Ensuring XYZ LAN Security	656
Appendix D. Notes on Electromagnetic Radiation	660
Appendix E. Export Restrictions on Encryption	662
Appendix F. Further Resources	664
References and Further Reading	664
Security-Related Organizations	665
Security Publications	667
Government Sources	670
The Cobb/NCSA Security Resource Disk	671
Appendix G. Online Glossary	672
Appendix H. How Public-Key Encryption Works	673
Modulo Math	673
Big-Time Math	674

Appendix I. Introduction to LANs	676
BN: Before Networks	676
The PC Arrives	677
The Network Idea	677
More Reasons to Connect	678
The LAN Arrives	679
Topologies	680
Network Protocols	681
NOS: Network Operating Software	682
 Appendix J. Securing Safe Software	 683
Legitimate Code	683
A Software Scenario	684
Software Sourcing	685
Further Software Precautions	686
 Appendix K. Appraising Microsoft AV	 689
The Antiviral Software of MS-DOS 6	689
Integrity Checking	695
Peculiarities in the Documentation	696
Bugs	697
Security Holes	697
Conclusions and Conjectures	700
Acknowledgments	701
 Index	 702
About the author	709

Acknowledgments

If you write a book about something, it is inevitable that people will refer to you as an "expert" in whatever that something is. However, writing a book about something also is a good way to learn how little you really know about that something. You also learn just how much you gain from the knowledge and wisdom of others, so I am pleased to have this opportunity to express my gratitude to the following, for sharing their knowledge and wisdom and thus making this book better than it would otherwise have been: Alan Fedeli, Alistair Philips, Ann Steffora, Bill Morone, Bill Zalud, Bob Bales, Bob Jones, Bruce M. Clay, Carol Ellison, Charles Cresson Wood, Charles Rutstein, Chris Goggans, Christopher Hughes, Christopher Klaus, Cliff Stoll, Craig Ellison, David Brake, David R. Johnson, David Kennedy, Derek Clifford, Don Parker, DT, Enrique Arroyo, Frank Derfler, Jr., Frank Ramos, Frank Andrew Stevenson, Fred Avolio, Gerry Faulks, Ira Winkler, J.D. Abolins, Jason Lamb, Jim Ross, Joel McNamara, John Podesta, John Taschek, Jonathan Wheat, Kore(sh), Laurence Milledge, Marcus Ranum, Martin Cheek, Max Schireson, Michael Miora, Otaku, Paul Brainard, Paul Constance, Paul Ducklin, Peter Gutmann, Paul Robinson, Peter Yueng, Ramon Barquin, Randy Nichols, Randy Pollack, Richard Charles Graves, Rob Rosenberger, Robert Steele, Roger Thompson, Ron Erdnman, Ross Greenberg, Sal Salamone, Sarah Cain, Sarah Gordon, Sylvia Moon, Tim Stefanini, Vincent Schiavone, William Hugh Murray, William Stallings, Winn Schwartau, Y. Radai, and not forgetting doctors, Alan Solomon, David Stang, Frederick Cohen, Mich Kabay, Peter Tippet, and Richard Ford.

With a list so long, there are bound to be people that I have omitted, for which I apologize in advance. Please let me know (scobb@ncsa.com), and I will remedy my sins of omission in the next printing and make sure that you get your copy of the book like everyone else. It goes without saying that, while these people contributed wisdom and knowledge, final responsibility for what is said between these covers, including anything unwise or erroneous, rests with me alone (please feel free to send me any corrections, or even disagreements, that you feel are appropriate—all that are included in future editions will be acknowledged and rewarded in the usual manner).

Putting any book together takes a lot of time, patience, and understanding from a lot of people. This one seemed to require more than most, and I am indebted to the following: Acquisitions Editor, Jenrufer Holt DiGiovanna, for keeping faith with the

xii Acknowledgments

project through so many interruptions; Senior Editor, John Baker, for pulling it all together; Stephen Moore and the rest of the Blue Ridge Summit crew for years of loyal service; Bob Bales, for seeing that Cobb/NCSA was a win-win combination; and of course, Chey Cobb and Erin Romfo, for their enduring love and support, and for being there, even when I wasn't.

Introduction

Security can be defined as freedom from risks and dangers. In this book, I show you how to minimize the risks and dangers of using personal computers. By personal computers, I mean everything from a global multiserver network to the PC in the den and the Macintosh in the classroom.

In recent years, we have come to rely on these machines to process and store vast amounts of increasingly important information. In many cases, the results have been positive. Impressive improvements in efficiency have been documented. There have been significant gains in personal productivity. Exciting new possibilities in education and entertainment have been enabled.

However, there are downsides to this new technology. Here are just a few of the current threats to the security of information entrusted to personal computers:

- Hackers, spies, and thieves

- Viruses and Trojan horses

- Malicious employees

- Disgruntled ex-employees

- Errors and omissions

In this book, I show you how to defend against these and other threats, such as fire, flood, and power failure. I also tell you how to guard against less obvious dangers, such as lawsuits over issues of privacy and copyright.

By reading this book, you will become familiar with the language and terminology of *infosec*, or information security. Perhaps most important of all, you **will** learn how to identify and defend against new threats, even ones which have not yet been created.

Why Choose This Book?

Many books have been written about computer security, data security, and information security. This book is different because, although it encompasses these subjects, it focuses on the machines that are on the front line of the information revolution.

that is, personal computers. Compared to most other computer security books, I think this book is more:

Realistic

Practical

■ Current

Enjoyable

Before I attempt to justify these claims, I should make it clear that this probably is not the only computer security book that you should buy. I honestly believe that it should be the first one that you read, but it is impossible for one volume to cover all aspects of this vast subject in enough depth to meet the needs of all readers. For this reason, I have provided plenty of references to other books that do a good job of addressing specific issues and technologies.

Realistic

This book shows you how to minimize the risks and dangers of using personal computers. It does not promise to eliminate them. Absolute security of information is entirely unrealistic. The best that we can hope for is *maximum* security, within the twin constraints of economics and utility. We want strong locks rather than doors that are welded shut.

While we cannot afford to ignore any threats that are even remotely possible, it is realistic to focus on those that are most probable. Being realistic also means being fully aware of just how devious and unpleasant some people who work with computers can be, without becoming paralyzed by predictions of gloom and doom. In reality, a lot of very clever men and women are working hard, and with considerable success, at keeping our information safe and sound. This fact tends to be obscured by the headlines, which irresponsible members of the media award all too often to people who are prepared to subvert technology to their own misguided ends.

In this book, you will find the facts calmly presented and the threats realistically described. The available responses are impartially assessed. I have no desire to frighten you into buying this book or any other security products or services. However, it would be irresponsible of me not to try and alert you to the problems that can befall those who are not sufficiently security-conscious.

Practical

I have endeavored to write a practical handbook rather than an academic treatise. The theory of information security has an unfortunate tendency to lose touch with the practicalities of real world implementation. In the real world, the costs of security measures cannot be ignored. They must be weighed against the value of what is being secured. Furthermore, because making information more secure often reduces its accessibility, potential security measures need to be discussed not only in terms of capital costs, but also in light of administrative feasibility and practicality, which includes both compliance and comfort levels.

It is a sad fact of business life that the budget for security, like that for employee training, often is inadequate. For this reason, I have included tips for "selling" security to upper management. I also have suggested low-cost solutions wherever possible, thus enabling you to get the most bang for each security buck. To help you find the tools that you need to create a more secure computing environment, the disk that is included with this book contains an extensive product directory, complete with vendor names and contact numbers. The directory is in hypertext format and can be browsed and searched in many different ways to help you find exactly what you are looking for.

Current

No book about computer security can claim to be completely current, even this one. New threats emerge on a daily basis, yet the production cycle of a book is measured in months rather than weeks. For that reason, even quarterly journals and monthly magazines have a hard time keeping current.

That is why each chapter of this book has its own page on the World Wide Web. The address for each page is listed at the beginning of the chapter. I have called this feature "W.E.B." for "Web-Enabled Book," and it will provide you with up-to-date sources for the very latest computer security news.

Enjoyable

It is unfortunate but true that both *computer* and *security* are words with a high *snooze* factor. Whether you are reading about them or listening to someone talk about them, it takes only a few sentences before eyes start glazing over. I am aware of this from my own experiences addressing seminars and conferences.

The most effective antidotes seem to be humor and pictures. For this reason, you will find that these pages contain some of the former and a fair number of the latter, in some cases attempting to save paper by combining the two, as in Figure 1.1. Beyond this, I have included several kinds of sidebars. These include Newsflashes, Flashbacks, and Pointers, all of which are marked by handy icons:



Indicates a security related news story, statistic, or case study.




Points to an example of changes in computer security since the first edition of this book appeared.



Draws attention to practical tips, tricks, or other valuable suggestions.

Is This Book for You?

Probably! This book is for everyone who uses a personal computer, plus anyone who manages people who use personal computers. That includes so-called SOHO users (small office/home office), power users, network administrators, and information technology managers. Even if you are a corporate executive who never touches a computer, this book is for you. One day, you will want to know what it means when



FigureInt.1 If possible, avoid access controls that are too complex or conspicuous.

the newspaper headline says "Virus Cripples Major Bank" or "Hackers Ransom Company Secrets."

To spread my message of constructive concern as far as possible, I have taken pains to avoid excessive use of *computerese* (the strange murnbo-jumbo mumbled by men with taped-together glasses and too many pens in their plastic pen protectors). When special terms are used, they are explained within the text. Acronyms are unavoidable in this business, but they have been kept to a minimum, and those that are used are explained.

How This Book Evolved

This book is based on an earlier work, *The Stephen Cobb Complete Guide to PC & LAN Security*. The initial performance of that book, which first appeared late in 1990, was disappointing. There were a few reviews, and they generally were positive, but sales never really took off. I began to wonder if anybody was really interested in personal computer security.

At that time, I was accustomed to the performance of software-specific titles, like *Using Quattro Pro*, sales of which are strongest when the book is first released but that, in many cases, recede to a trickle shortly thereafter. However, while the *Guide to PC & LAN Security* started out moderately, it stayed the course, continuing to sell at a steady pace even as other books that I had written at about the same time were remaindered.

This fact came to my attention about the same time that I received a number of invitations to speak about various aspects of computer security. It began to occur to me that concern about security issues finally was starting to spread beyond the in-

ner core of specialized security professionals, a pleasantly surprising number of whom, I have since discovered, actually owned a copy of my book.

In the summer of 1994, I received confirmation of this trend when I joined the NCSA Information Security Forum on CompuServe. There I found a veritable hive of activity. Thousands of people were asking and answering questions and uploading and downloading security-related files, from data protection programs to hacker alerts and news bulletins.

Soon I was joining in, and it was not long after that I became a sysop (system operator, a person who answers questions, moderates discussions, keeps an eye on things). That was the start of my relationship with the NCSA, and since then, we have worked on several projects together. I continue to sysop in the PC/MAC/LAN Security section, although I am active in several other areas as well.

Although the *Guide to PC & LAN Security* was proving quite durable by computer book standards, by the end of 1994, some sections were looking dated. The predicted growth in networking and telecommunications had happened, and there was a lot more that needed to be said about these subjects. Hacking had taken an increasingly ugly turn. Clearly a new book was needed, and I sat down to write it, encouraged, assisted, and eventually employed by NCSA. It took over a year's worth of effort, plus a lot of patience and forbearance from editors Jennifer Holt DiGiovanna and John C. Baker, as well as the unwavering support of friends and family, but the volume that you have in your hands is the not insubstantial result.

At the NCSA, I discovered a widely respected team of people dedicated to promoting safe, ethical computing. I determined that, if possible, I would not simply revise my *Guide to PC & LAN Security*. I would relaunch it under the NCSA banner. However, the motivation for this new edition remains the same as for the first.

This book was written because many users and managers of personal computer technology are worried about risks and dangers. Obviously, some risks and dangers are very real, but some are imagined, the product of an oversensational, underinformed media.

Effective defense against some risks and dangers is difficult, but many can be thwarted by simple precautions. A very real concern is that the fears and worries of management and end-users will inhibit the spread of personal computing's positive benefits.

The mission of this book is to restore some of the sense of security and spirit of optimism that has marked previous stages in the growth of personal computing. If your comfort level with the systems that you use or manage is increased by what you read here, then that mission is being accomplished.

Where to Find What

Although this book works best if you read it from cover to cover, here is a brief outline of the contents in case you have a specific and immediate area of interest:

The first three chapters describe the problem of personal computer and network security, then set out some straightforward and immediate security measures. Methods of risk assessment and security planning are discussed.

Chapters 4 and 5 focus on securing and supporting hardware, including power supplies, theft deterrence, and physical restraints.

Access controls are addressed in chapters 6 and 7, which cover site access, system access, file access, and password schemes. Techniques for controlling removable media also are discussed.

Backup and recovery are covered in chapter 8, which presents strategies for reclaiming data and recovering from disasters such as fire and flood.

In chapter 9, the computer virus threat is discussed and antivirus measures are outlined. Tips for dealing with a virus attack are presented.

Chapter 10 highlights a major source of virus infection, pirated software, which also can expose your organization to hefty fines and court cases. Measures to identify and reduce your exposure are detailed.

In chapters 11 through 13, the focus is networks and communications and their security implications. Remote access, firewalls, and Internet connections are examined.

Questions of ethics and motivation are discussed in chapter 14, which takes a closer look at hackers and other human threat categories.

Chapter 15 provides a review of the main points in each chapter and draws some conclusions for present and future information security practice.

Special Foreword by the Author

Welcome to the tenth anniversary edition of *The Stephen Cobb Guide to PC and LAN Security*. I trust that these pages will provide you with a solid yet readable introduction to computer security.

Most of the text in this book first appeared in 1996 as *The NCSA Guide to PC and LAN Security* (McGraw-Hill). That book was a revised version of the original *Stephen Cobb Complete Book of PC and LAN Security*. I started writing that book in the late eighties for TAB Windcrest, which was later acquired by McGraw-Hill.

Something Old, Something New

The two previous versions of this work have been read by hundreds of thousands of people. Many colleges and universities have used them as textbooks. For a long time they have been on two very different reading lists: the alt.2600 FAQ and the ISC² CISSP Examination Study Guide (I will say more later about what both of those mean). This created a steady demand for the book. Supplies eventually sold out. After numerous inquiries from people who could not find the book at book stores, I decided to take matters into my own hands. I took the necessary steps to republish it. I also chose a less expensive format, and made it sure it would be available around the world, through leading online booksellers.

Reprinting was a chance to make a few changes. The previous version included a floppy disk. This increased the cover price. I decided to drop the disk. This helped lower the price. Instead of being on a disk, the files that came with the book are now online, accessible via the World Wide Web. The following web page is your starting point for these files and other useful information, some of which is only available to buyers of this book:

<http://www.cobb.com/pclan>

This location will also provide you with more about Internet security than you will find in the current version of the book. However, if you are new to computer security think twice about tackling the subject of Internet security until you understand the basics of personal computer and network security, as outlined in the book.

The Internet is clearly the biggest single change that has taken place in the field of computers and computer networks during the last ten years. But a surprising number of threats to computer security are essentially the same today as they were back then. The same is true of many of the defensive measures taken to protect the confidentiality, integrity, and availability of computer-based information.

The Human Factor

Why are a lot of security issues the same today as they were ten years ago? Computer security is more of a people problem than it is a technology problem. Sure, there are technology issues. For example, computer systems need a clean and constant supply of electrical power (which you can read about in Chapter 5). But apart from power spikes and outages, fire, flood and other natural disasters, the biggest threat to computer-based information is human nature. If human beings did not have a tendency to pry, lie, cheat and steal, we wouldn't need things like passwords, encryption, locks, anti-virus programs, and so on.

The fact is, much of the work of computer security professionals involves countering the effects of such common human traits as laziness, and forgetfulness. We're talking about things like leaving backup tapes lying around, or not scanning new files for computer viruses, or using weak passwords (described in Chapter 7—weak passwords are ones that are easy to remember and thus easy for attackers to guess). A lot of computer security is common sense and sound business practice, but just because we humans know what we should do, that does not mean we always do it. Long before the Internet became a business tool, many network users had developed bad habits. Indeed, in computer security terms, the most important role of the Internet so far may be to hold up a mirror to the accumulated security problems that arose earlier. I am referring to the evolution from mainframes to networks of personal computers (upon which most of the world's computing is now performed).

Which brings me to a couple of references I made earlier: 2600 and ISC². These are two different organizations which provide two very different views of information technology, although they share a penchant for numeric names. The name ISC² refers

to the IISCCC, or International Information System Security Certification Consortium. This is the non-profit body which administers the CISSP certification (Certified Information System Security Professional). The CISSP is the leading, non-commercial, professional certification for people working in the field of computer security.

To be precise, the CISSP covers information system security, of which computer security is a sub-set (information system security is itself a subset of information security, which includes non-digital information, such as paper documents). You can find out more about CISSP at www.isc2.org. Becoming a CISSP involves more than just passing an exam. You have to have field experience and must subscribe to code of ethics. You must also maintain your professional knowledge or you lose your certification.

The CISSP exam is based on something called the Common Body of Knowledge or CBK. This is defined as what an experienced information system security professional should be expected to know. The book you are holding is considered to be part of the CBK. This does not mean that reading this book will teach you everything you need to know to pass the CISSP exam. But reading it won't hurt either.

A very different entity from ISC² is 2600, publisher of "2600: the Hacker Quarterly." You can find this publication on magazine racks in some countries. It is entirely supported by sales and subscriptions, with no advertising. The content includes tales of hacking and tips on how to hack, mainly within the classic definition of that term, that is, the non-malicious investigation, and tweaking, of technology.

It is important to remember that back in the seventies people like Steve Jobs and Stephen Wozniak, the founders of Apple Computer, referred to themselves as hackers. Indeed, before they made computers, Jobs and Wozniak made devices called Blue Boxes, and sold them door-to-door on campus. They learned how to make Blue Boxes, which are capable of gaining free access to the phone network, in violation of the law, from a man called John Draper, a.k.a. Captain Crunch.

This "handle" came from the fact that, in the early 1970's, Captain Crunch cereal was sold with a small whistle in each box, and John Draper was one of the first people to realize that the whistle generated a tone or signaling frequency used in long distance telephone lines. The effect of injecting this signal into a line was: "to cause the remote long-distance exchange to terminate the connection and listen for touch-tone signals encoding the new destination being called...the billing for the long-distance call ended with the whistle tone, and the new long-distance call was made at no charge to the customer" (Douglas W. Jones, University of Iowa Department of Computer Science).

The tone of the Captain Crunch whistle was 2600 cycles, hence Draper's handle and the name of the hacker newsletter, early editions of which predate the Internet. They circulated on bulletin boards in a very basic electronic format. The first version of my book was listed in an FAQ, a list of frequently asked questions, which the newsgroup alt.2600 published. It was referenced under the heading: What books are available on this subject? It was not specifically recommended and there was no suggestion that the book was a guide to hacking.

I'm just letting you know that if you search the web today for "Cobb, hacking, 2600" you will still see it listed. So bear in mind that hacking, at least in its early or "pure" form is not necessarily a criminal activity. Some people will use the term "cracker" for people who commit criminal acts using skills associated with hacking. But I have a problem with that since the word cracker is already in use to identify a different kind of person altogether, as anyone who has lived in The South can attest. A good friend of mine, Dr. Mich Kabay, advocates the terms "criminal hacker" and "criminal hacking" to distinguish between people who respect the law when using hacking skills and those who do not. Understanding this distinction is an important part of becoming a computer security professional.

A Word of Thanks

I would like to close this foreword with a word of thanks to all of the people who have encouraged and contributed to my computer security efforts. These range from my immediate family to people I have only met electronically: Chey Cobb, fellow CISSP and perfect partner, has inspired me with her struggle to protect our country's digital assets (check out www.cheycobb.com); brother Mike Cobb, MCDBA and founder of Cobweb Applications, Ltd., has kept me enthused and taught me a great deal about the bits and business of security software (see www.cobweb.co.uk); mother Dorothy Cobb, a tireless force for human progress, has performed Herculean tasks of editing and organization for me; the "lab guys" have been an enormous source of wit and wisdom, illuminating many areas previously unknown to me (thanks Don, Ryan, Devlin); the "infosec guys" have kept me on my toes and prevented me from taking myself too seriously, as well as providing wise counsel and truly excellent security thinking (thanks Michael, David, Vince, and Bernie). Others in the field who have educated and inspired over the years include Mich, Winn, Richard, Sarah, topher, and others too numerous or covert to mention. And of course Ron Powers, who agreed,

back in 1989, to let me submit this book for publication instead of a manual for an application which never took off. My sincere thanks to one and all.

Stephen Cobb
Fairfax, Virginia
March, 2001

Security Matters

Assessing the Problems, Threats, and Issues

*"The ~~time~~ is 9 A M Do you know where your
data is?"*

This chapter maps out the subject matter of the book, defining terms and assessing the current state of personal computer security. There are several self-tests that you can perform to measure your personal computer security awareness. There also are some "myth-busters," which might help to reassure readers who have been panicked by OTT press reports (OTT as in Over The Top, not a strange new computer acronym). While you already might have a high level of security awareness and a strong urge to spring into action, please begin with this chapter. It will help you place your actions in a broader context, thus making them more effective in the long term.

The State We Are In

Definitions can be pretty boring so, before we get into a discussion of what security means to computer users, we will consider some facts and figures that put the subject in perspective.

War stories

Many readers probably have war stories of their own. Indeed, you might have picked up this book in response to a security incident involving you and/or your organization. If you have been lucky enough to avoid any problems so far or happen to have

been avoiding television, radio, newspapers, and every other media for the last five years, the following sample of facts and figures will give you some indication of what is happening:

- An estimated \$1 million worth of computer equipment, believed stolen this summer from a terminal at Logan Airport, has been found in an East Boston garage where authorities also discovered 18 pounds of marijuana in the trunk of a 1967 Lincoln Continental. Officials said the raid was part of an ongoing investigation by Massachusetts and New Hampshire State Police into a computer theft and drug distribution ring. (Boston Globe, September, 1994)

The total number of personal computers in use today, worldwide: 100 million. (Personal estimate, March, 1995)

Percentage of survey respondents running "mission-critical" applications on local area networks: 76%. (Ernst & Young survey of 1271 technology and business executives, January, 1995)

Unemployed 26-year-old Christopher Pile appeared in court on February 14th accused of writing and distributing viruses. Pile is the first person to appear in court in the U.K. to face charges relating to computer viruses under the Computer Misuse Act. Pile is accused of writing the SMEG viruses, Queeg and Pathogen. Each offense carries a maximum penalty of five years in jail. (Secure Computing, March, 1995)




Figure 1.1 A representation of the worldwide trend in computer-related crime, fraud, and abuse

The chart shown in Figure 1.1, which was based on numbers published by the Royal Canadian Mounted Police, is representative of the worldwide trend in computer-related crime, fraud, and abuse.

Barely more than half of respondents actually had tested their business resumption/recovery plan in the last 12 months. More than 1 in 5 did not have a plan. Results of the annual industry survey, performed and reported by Infosecurity News, March, 1995

Of 9000 computers attacked, 7900 were broken into. Only 4% of the successful attacks were detected by the target organization. Of those 320 organizations, only 5% actually reacted to the attack. Report in Washington Technology describing tests performed by the Defense Information Systems Agency against Department of Defense computer systems, 1995.

Only 22% of respondents said that management viewed security as "extremely important," which is way down from last year's 34%. Ernst & Young survey of 1271 technology and business executives, January, 1995

The year in which, at current rates of growth, there is projected to be an Internet connection for everyone on the planet: 2009. Input, Inc., Mountain View, California

- The Computer Emergency Response Team reports a 77% increase in computer break-ins in 1994 over 1993. Infosecurity News

American Airlines is suing Northwest for \$50 million over the alleged theft of confidential data, mailed on a floppy disk to Northwest's head office in Minneapolis by an American Airlines employee, who went to work for Northwest shortly thereafter. Report posted to the NCSA Information Security Forum, 1994

The number of high density 3.5" diskettes manufactured in 1994: 4 billion. Avanti Associates, Valley Forge, Pennsylvania

- More than 40% of respondents estimated that losses would exceed \$5 million if their organization's computer data was tampered with, erased, lost, or stolen. Annual survey, Infosecurity News, March, 1995

Increase in the number of reported incidents of fraud and abuse, between 1990 and 1993: 300%. Increase in the total cost of reported incidents: 183%. Report on the fifth triennial survey of the extent of computer fraud and abuse, Audit Commission, U.K.

Between 40% and 68% of survey respondents felt that security for mission-critical applications was inadequate, depending on the platform. Ernst & Young survey of 1271 technology and business executives, January, 1995

What's your problem?

The previous incidents and statistics give you some indication of the current threats to information security. However, which are the most important threats? Which should you be most concerned about? I cannot give a categorical answer to this question for two reasons. First of all, the answer varies from user to user and system to

system. If you are installing a network for a bank whose employees have little previous experience of computers, then the answer probably is errors *and* omissions, which is a catch-all phrase for unintentional losses caused by humans. However, if your organization has been using computers for some time and now is experiencing layoffs and cutbacks or has a high level of staff turnover, you might be more at risk from disgruntled and dishonest employees.

The second reason why it is hard to say which threats are the most prevalent is that a lot of security problems are never publicly discussed, for obvious reasons. Suppose a bank employee has found a loophole in the computerized accounting system that allows her to divert other people's money to a friend's account. This is not something that the bank wants to make public because, however wrong the employee might be, the bank still ends up looking less than reliable.

What I can do is give you a "best guess" based on the accounts of incidents that are made public plus the combined experience of a number of security experts. The chart in Figure 1.2 is used in security awareness presentations by the NCSA and reflects input and opinion from a wide range of sources.



Mission Critical


I have used the term mission critical several times now and, despite the fact that it sounds a bit like military-speak, it is a very handy, and fairly self-explanatory, piece of jargon. It typically is applied to any hardware, software, process, task, or other component of an organization's mission that is critical to the successful completion of that mission. The term also is conveniently relative. We can say that an airline's computerized reservation system is a mission-critical application, as is the word processing program that I am using to write this book (if the program screws up and scrambles my text, the book is in jeopardy).

A typical case?

The unintentional errors that people commit probably are the biggest single cause of data loss. If we group together the people who are dishonest, disgruntled, hacking, or writing viruses, then the next largest threat also is people. However, not all security problems come down to folly, fate, and felons. Some arise from the realities of everyday office life.

Consider the following case notes and ask yourself if they sound familiar (don't worry if some of the terms are unfamiliar, they will be explained later). During the fourth quarter of 1994, an associate was contracted to perform a physical audit of computer equipment at the regional offices of a high-tech company. This company sells goods and services to both businesses and individuals and has a fairly high profile in a very competitive market.

The office occupies five floors of a glass and concrete building just off the Interstate. The company estimated that it had about 160 computers, both IBM-compatibles and Apple Macintosh, at this site. Most were connected to a local area network (LAN) running Novell NetWare. This was connected to a wide area network (WAN) with other regional offices in the state.



Errors and Omissions

Figure 1.2 A breakdown of the "best guess" as to security threats, per NCSA.

Here's what turned up in the way of undocumented hardware — some of it in use, some of it packed away in closets:

- 12 XT/AT machines
- 20 486 machines
- 3 dozen printers
- 2 dozen monitors
- 6 modems

In other words, the company owned 20% more computers than it realized. This sounds better than finding 20% of your hardware is missing, but it represents poor facilities management, shows a lack of budgetary control, and is symptomatic of weak management in other areas.



Modems Verboten?

"Prohibit modems on PCs that you do not know about, and periodically check for illegal modems." Daniel E. White, national director of Information Security Effectiveness, Ernst & Young, as quoted in *Information Week*, January, 1995.

If an organization cannot keep track of how many computers it has, you probably would predict that the control of information stored on those computers to be less than adequate. You would be right. Without using any special hacking tools, our associate gained access to 70% of all the personal computers and 30% of all the network connections (in this context, the term hacking is used in the narrow sense of "overcoming obstacles to computer access").

Although many of the systems were password-protected, she found the passwords were easy to guess or clearly visible. A common practice was to write the password on a piece of paper stuck to the monitor. One user obviously found this too informal and recorded his password in embossed tape, which he stuck to the bottom of his keyboard. Many passwords were either first name, last name, or the company name. Clearly the company did not enforce any sort of password regime.

Only a few of the computers had BIOS passwords enabled to prevent system access. On most system units that had case locks, the keys still were in the lock (see Figure 1.3). Keyboard locks were not in use, and it was common to find these keys still in the machine. There was no central repository for keys. Access to case lock keys not only helps people steal the contents of the system unit, such as memory chips, it also allows them to get around BIOS passwords. Conversely, if a person set a BIOS password, then left the company without telling anyone what it was, either the case lock or the password would have to be broken to gain access to the system.



Information Overload

One of the biggest challenges when writing about computer security problems today is knowing when to stop. When I was putting together the first version of this book in the late 1980s, it was not that easy to find examples of security snafus or case studies of security breaches. Now I have to ration myself or my filing cabinet will collapse.

Another major security weakness at this office was a lack of adequate access control. Members of the public visited the office for both customer service and sales but were not properly confined to specific areas. Anyone wearing a suit and carrying a




Figure 1.3 Circular pin tumbler keys of the kind commonly used to lock out keyboard input and prevent system cases from unauthorized access.

clipboard or file folder could walk into any area unchallenged and take the elevator to any floor. The one exception was the floor reserved for accounting. Elevator access to that floor was controlled by a coded keypad. However, there was nothing to prevent you from riding the elevator to the floor above, then taking the stairs back down to that "restricted floor."

Note that, in this example, as in so many cases, breaching the security of computerized data requires no special knowledge of computers. Competitors of this particular company would have had a field day with nothing more than the ability to read. That's because numerous printed reports were left in unattended printer output trays for considerable lengths of time. These include lists of customer names, account numbers, and balances. One associate particularly enjoyed reading the contents of an unattended wastebasket clearly labeled "Documents for Shredding." If someone was going to look through the trash cans in this office for "useful" information, this would be the obvious place to start.



Dumpster Diving

One of the classic tools of the hacker is dumpster *diving*. This involves looking through a company's trash, which often is found awaiting collection in dumpsters behind the building, for documents that provide valuable inside information. Much of the fraudulent use of telephone lines that plagues phone companies was made possible by people who retrieved discarded phone company manuals and schematics from the trash.

The point of this case study of a "typical" office is not to point fingers or poke fun. There are many offices around the world that are just like this one. They represent accidents about to happen, lawsuits about to be filed, data about to be stolen, money about to be lost, and operations about to grind to a halt. We need to raise security awareness and put an end to this way of doing business.



Read All about It

A good place to catch a roundup of the latest stories involving computer security incidents is the News/Case Studies section of the NCSA Information Security Forum on CompuServe (see appendix F for details). The News/Case studies library contains a comprehensive archive of security reports.

Now the Good News

At the beginning of each episode of the TV show "The Bionic Man," they used this phrase: "We have the technology!" That sums up much of the good news for people interested in protecting computer-based information. In other words, there now are plenty of products on the market to help you plan, design, implement, and monitor computer security measures.

We have the technology

Each year, the magazine *Infosecurity News*, working in conjunction with the NCSA, assembles a Buyers Guide, listing information-security products. The 1995 edition lists nearly 1000 products, a 15% increase over the previous year. As you can see from Figure 1.4, the database is available on disk as a hypertext catalog.

There now are more than 400 companies that sell tools with which to analyze, develop, implement, and maintain security procedures. These encompass both hardware and software solutions. The hardware ranges from simple cable and padlock systems to deter theft of computers, printers, and other equipment to sophisticated encryption and access control systems. In Figure 1.5, you can see an example of the latter: the SecurID Card.

The purpose of the SecurID Card, produced by Massachusetts-based Security Dynamics, is to give the user a new password every 60 seconds. This appears in a small LCD panel on the card. When the user wants access to the secured system, he or she must enter the code that appears on the card. A clock within the card is synchronized with the clock on the system being accessed so that the password will be recognized. However, because the password is only good for 60 seconds and a new one is presented to the user whenever one is needed, there is no need to write it down and thus compromise it.



What Is Encryption?

Encryption is a cryptographic term. "*Cryptography* is the art of creating and using cryptosystems." This is the definition provided by the cryptographic experts who wrote the excellent FAQ put out by the *sci.crypt* discussion group on the Internet (an FAQ is a set of answers to Frequently Asked Questions). "A *cryptosystem* or *cipher system* is a method of disguising messages so that only certain people can see through the disguise." The original message is called a *plaintext*, and the disguised message is called a *ciphertext*. "*Encryption* means any procedure to convert plaintext into ciphertext. *Decryption* means any procedure to convert ciphertext into plaintext."

In the context of computer security, encryption is what happens when you password-protect data, such as a record within a file, a complete file, or an entire disk. The art of breaking or cracking cryptosystems, that is "seeing through the disguise even when you're not supposed to be able to," is called *cryptanalysis*. The study of both cryptography and cryptanalysis is called *cryptology*. There will be more about this in chapter 7.

Software for security work is available in abundance. There are programs to analyze security risks, audit all forms of computer activity, and to password protect everything from entire systems to individual files. One of the most powerful forms of encryption, the RSA algorithm, now is available in commercial products, such as RSA Secure, which is shown in Figure 1.6.

RSA Data Security, which creates and licenses products based on the Rivest-Shamir-Adleman algorithm, announced its RSA Secure hard disk encryption soft-

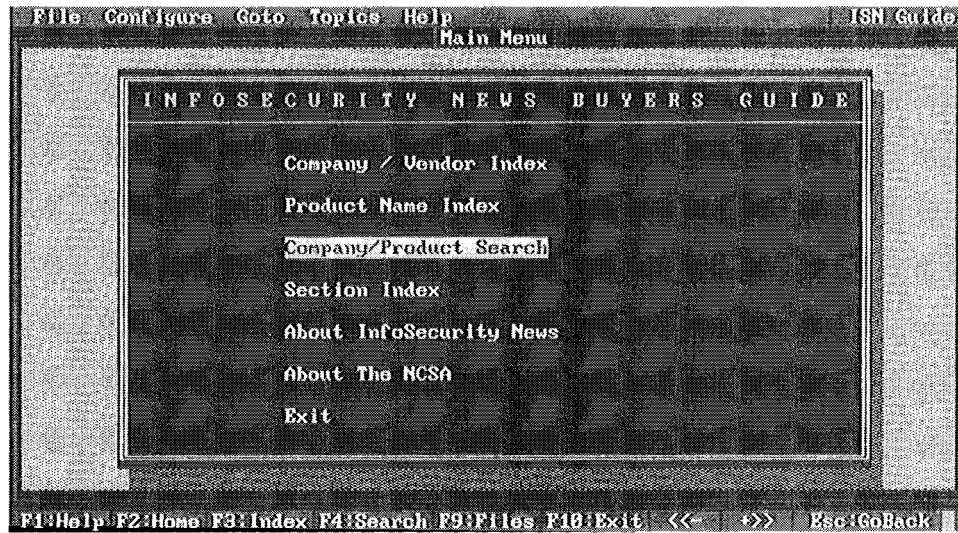


Figure 1.4 The *Infosecurity News* database, which can be found on the Cobb/NCSA Security Resource Disk included with this book




Figure 1.5 SecurID, a credit card-size token for producing one-time passwords

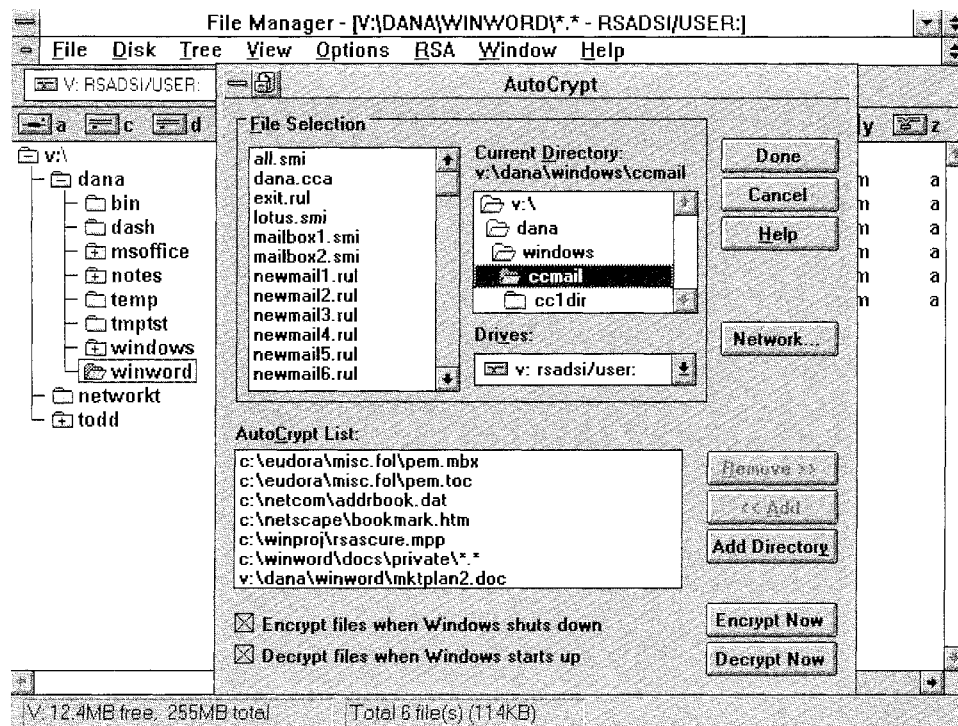


Figure 1.6 RSA Secure is a powerful hard disk encryption program

ware for Windows in January of 1995. It is particularly useful for firms with fleets of remote PCs and includes the ability to recover files created by workers who are no longer with firms and decrypt the files.

Turning the tide

Are you fed up with the high interest rates and burdensome service charges on your credit cards? Well, one reason for those charges are people who use their computers to break into other people's computers and steal lists of credit card numbers. These then can be used to fraudulently charge phone calls, goods, and services in the names of innocent victims. One person who did steal credit card numbers, but denies using them, is Kevin Mitnik.

If we are lucky and play our cards right, February of 1995 will mark the turning point for information security experts. This was the month in which the fugitive Mitnik was apprehended by the FBI, largely due to the efforts of Tsutomu Shimomura, a physicist who lives in California. In recent years, the FBI has had some success finding and prosecuting hackers who commit fraud, but what was different about this case was the decision of some news media, such as the *New York Times*, to highlight the good guys, in this case Tsutomu Shimomura and his associates.

There are several theories as to why the American media has tended to glorify hackers, which is an issue addressed in detail in chapter 14, but there is a parallel tra-

dition of revering super-heroes, and it is to be hoped that, in the future, more of the people who toil to keep our systems safe and sound will be getting the praise and encouragement that they deserve. People like Cliff Stoll, the Berkeley astronomer whose curiosity over a 75¢ discrepancy in a bill for computer time led to the tracking of a hacker all the way to a KGB spy ring. You can read the full story in his excellent book *The Cuckoo's Egg* or catch the cyber version on the Internet (see Figure 1.7).



Hacking beyond the Pale

A suspect in a credit card scam was arrested in mid-January, leading to a full-scale investigation into a national child pornography ring that uses stolen computers to fund the project. Police said that Darren Balzarini was arrested after a tip-off from Federal Express and was found to be in possession of \$250,000 of computer equipment, allegedly purchased with stolen or forged credit cards. The equipment was to be used to finance and effect the distribution of child pornography across the U.S. using a BBS-style system.

Secure Computing, March, 1995

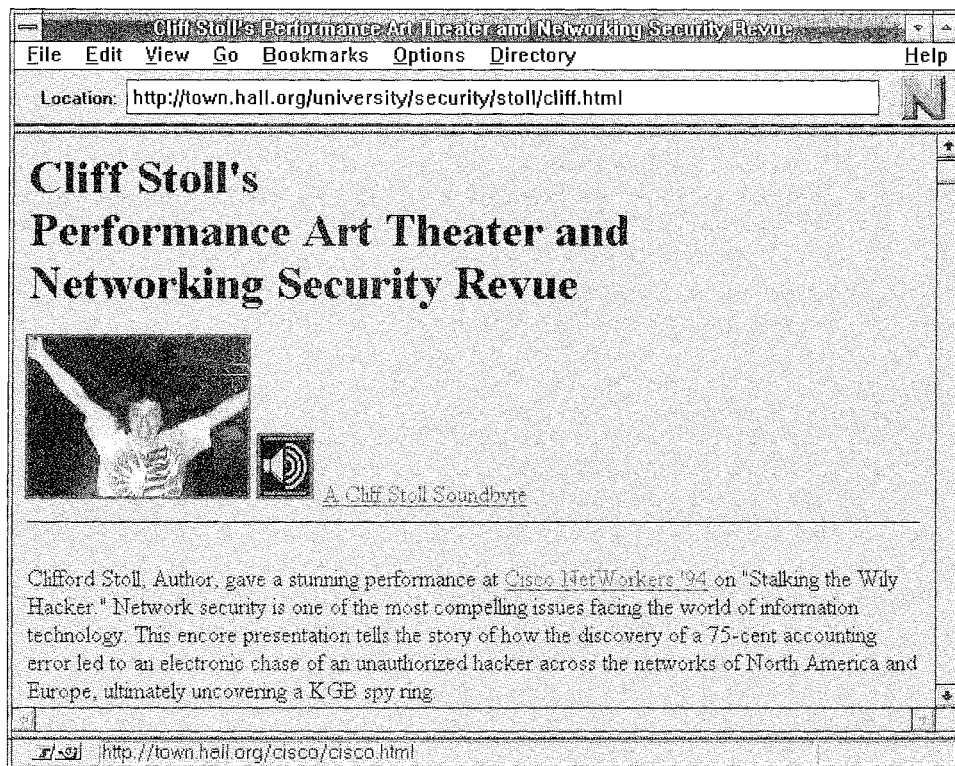


Figure 1.7 Excerpts from Cliff Stoll's amazing story can be found on the World Wide Web, complete with sound bites

Defining Security

The Introduction to this book defined security as freedom from risks and dangers, but we need a more specific definition of information security. The classic definition is easy to remember by its acronym, CIA, which stands for *confidentiality, integrity, and availability*.

Confidentiality

If a computer security system is effective, disclosure of information handled by that system will be limited in strict accordance with the wishes of the owner of that information. The term "owner" means the person who created the information or holds the copyright to that information. This is an important point because a lot of hackers are fond of saying that information belongs to everyone or "information wants to be free." To such people, the "ownership" of information is somehow anathema.

However, if I write a letter with my word processing program, it sure as heck belongs to me, and I have every right to control to whom the contents are disclosed. Likewise, a company that develops a new and improved process for painting cars is entitled to allow or deny access to that information, regardless of whether or not it is stored on or transmitted between computers. Thus security systems attempt to guarantee to the owners of information that their disclosure will be effectively controlled.

Integrity

Keeping the wrong people away from my information isn't just about keeping things secret, it also is about ensuring that my information remains complete, whole, intact, and uncorrupted. I need to know that, when I print out the first quarter sales figures, they will not have been altered, either accidentally or on purpose, since I entered them. Likewise, if I am using a network of personal computers to pay invoices, I need to be sure that all of the invoices that are paid are legitimate. Security systems attempt to guarantee the integrity of information and guard it against all manner of threats, from people to power surges.

Availability

In its 1994 survey of management and technology executives, Ernst & Young found that 32% reported experiencing "nonquantifiable losses" due to a lack of system availability. In other words, the information stored on their computers was not available when they needed it. The processing capability that their organizations' computers normally provided was not available when it should have been. Whether you look at it as maximizing "uptime" or minimizing "downtime," security systems need to deliver the data and the processing power whenever it is needed and provide the best possible recovery when disaster does strike.

Computer Types

Because the focus of this book is personal computer and network security, I need to define what these terms mean. All computer systems have four common elements:

input, processing, storage, and output (see the diagram in Figure 1.8). This activity can take place in a single box, in a single location, or be dispersed over a wide area. This distinction is one of the ways in which we categorize computer systems. Although the distinctions between different types of computer systems have been breaking down in recent years, it still is possible to roughly delineate three different categories.

Mainframes

The mainframe computer represents the original architecture of computing. Traditionally the mainframe is a large piece of equipment, consisting of racks of processing and storage equipment, stored in a purpose-built room with special climate and access controls. The mainframe supports multiple users via multiple stations, or terminals, as illustrated in Figure 1.9.

Because many of the first computers were developed to perform military and financial functions, the means to restrict access have always been inherent in mainframe design. Mainframe operating systems have extensive facilities for data protection, and full backup facilities are an integral part of any mainframe installation.

The terminal connection. To offset the high initial cost of mainframe hardware, it was put to use around the clock. This led to the need to provide access to users at locations other than the central processing unit itself. Access to a mainframe usually is by means of a terminal. A *terminal* traditionally is defined as an input/output device with no computational ability or storage capacity.

The terminology of personal computing owes a lot to the original mainframe architecture. For example, the term *workstation* still is used to refer to a physical location at which computing tasks are performed, whether they are performed on a mainframe terminal or a personal computer. However, the mainframe computer generally is assumed to be losing ground to less centralized systems, such as networks made up of extremely powerful personal computers.




Figure 1.8 The four main elements of a computer system.




Figure 1.9 Diagram of a mainframe computer system

The "mainframe mentality" of centralized command and control certainly has lost its appeal, replaced by concepts such as downsizing and client-server computing, despite the fact that MIS staffs still are struggling to define these concepts and understand their implications for security. (MIS is an acronym for management information systems.)



MIS Directions

You can get some idea of current trends in MIS from a survey published in 1995 by the Association for Systems Management, which looked at MIS-related job advertisements published in major U.S. newspapers. The job market for MIS services was shown to have grown by 44% during the 1993–94 period. Topping the list of skills most in demand in the MIS field were those relating to networking, Unix, personal computing, relational databases, Windows, and C or object-oriented languages.

But not dead yet. By the end of the 1980s, the logarithmic leaps in microprocessor performance led some analysts to predict the imminent death of mainframe computers. A *Computer Intelligence* study in 1992 indicated that 88% of IBM new ES/9000 mainframes were being sold to the existing installed base rather than to new customers. The same study suggested that sales of IBM and compatible mainframes had declined 11% in 1991.

In 1992, the industry journal *Datamation* reported that the percentage of new applications destined for mainframe was 13% against 87% for nonmainframe platforms. This was a significant shift from the 20180 breakdown just a year before. An International Data Corp study in 1993 saw the value of worldwide mainframe sales shrinking 0.4% per year from a 1991 peak of \$28.1 billion.

Despite these figures, it could be argued that mainframes will continue to handle more of the world's total data processing workload than any other technology for at least the rest of the century. That is because mainframes are firmly established at the heart of heavily transaction-dependent businesses, like banking and government (particularly the military and intelligence portions thereof). A decline in new business does not necessarily presage the inevitable demise of current mainframe activity.

The falling dollar value of sales is partly offset by declining prices. While an extensive Ernst & Young survey of technology and business executives at the end of 1994 found that mainframe use was decreasing at 36% of the organizations surveyed, the same study showed it increasing at 50% of them. So, while the mainframe might no longer dominate computing in the 1990s, it still plays a vital role in many organizations.

Minis

After the mainframe came the minicomputer, which was smaller yet still able to support multiple users. Minicomputers, such as the DEC VAX systems, found a place serving the entire computing needs of small organizations or most of the needs of departments within larger organizations. The minicomputer offers central storage and processing of information entered through inexpensive terminals, as illustrated in Figure 1.10. The minicomputer retained many of the security features of mainframe systems, such as sign-in procedures when terminals are used, automated backup of stored information, passwords to control file access, and locks on the terminals themselves.

In some organizations, the mini, which typically runs some form of Unix, or a proprietary OS, such as DEC's VMS, continues to play a valuable role. A survey by Ernst & Young in 1994 found that 72% of organizations were using minicomputers to run mission-critical applications. Some minis survived by being upgraded and turned into departmental servers, creating a connection between the mainframe and the network. However, in many cases, minis have been replaced by networks of personal computers.

Micros

Technically defined as a small computer designed around a central processing unit that is contained in a single integrated circuit or microchip, the microcomputer con-




Figure 1.10 Diagram of a minicomputer.

tains all of the basic features of a larger system, but together in a single unit serving a single user, as illustrated in Figure 1.11. The unit actually might be several small boxes cabled together, but the whole thing will fit on or under a desk, even into a coat pocket, providing one person with the ability to enter, process, store and retrieve data, without reference to or assistance from any other facilities.

The term *personal* computer is applied to microcomputers because microcomputers allow an individual to take charge of the entire computing process. None of the components have to be shared with other people for the system to work. Personal patterns and styles of work can flourish without the restrictions of conformity to central control. The situation is not unlike the automobile versus the bus. The personal computer provides freedom for the user to choose the destination, the route, and the schedule. With your own computer, there is no question of waiting for other users to log off or waiting for the system administrator to install the right software. You are the system administrator.

LANs, WANs, and Other Permutations

A personal computer does not need other computers or other resources to function. However, in some situations, it is advantageous to share or pool resources. This gives

rise to groups of personal computers connected, or networked, for the purpose of sharing. On the other hand, mainframes still have impressive capabilities, such as processing vast numbers of transactions in real time (meaning "as they happen"). Thus there often is a need for networks of personal computers to connect with mainframes to access this information.

Connecting a network to a mainframe can result in huge amounts of data traffic data. Sometimes it makes sense to use an intermediary, which can lead to any number of three-level arrangements, one of which is diagrammed in Figure 1.12. In the next few sections, I'll look at the various elements in this structure.

Terminals

A terminal is an input/output device for a multiuser system. A typical terminal is little more than a screen and a keyboard. It contains no storage or processing capabilities, hence the term dumb terminal. A personal computer can be used to emulate a terminal. When acting as a terminal, a personal computer does no real processing of data; it merely sends it or receives it. However, the personal computer can be switched out of terminal mode and get to work on processing data that it has received from the mainframe.

LANs

A local area network, or LAN, is a group of personal computers connected together to share resources. A good place to begin the LAN story is with storage facilities. The primary means of data storage on personal computers is magnetic media. Data and




Figure 1.11 Diagram of a personal computer.




Figure 1.12 A three-level computing structure using personal computers, a minicomputer/departmental server, and a mainframe.

programs are stored magnetically much like music is stored magnetically on a cassette tape and pictures are recorded on VCR tapes. The first microcomputers actually used audio cassette tapes to store programs and data. Cheap portable cassette players could be used to write information onto tape and then read it back.

Tapes were followed by *floppy* disks, so-called because they consist of a circle of flexible plastic coated with the same stuff that coats cassette tapes. Disks offer the advantage of faster operation. A read/write head can quickly be moved to any part of the disk through a combination of disk rotation and the movement of the head across the radius of the disk. Faster still are *hard* disks, which are sealed units that contain multiple magnetic coated metal platters. Precision construction allows the data to be packed tighter on the disk, providing greater capacity. It also means higher costs.

When the first hard disk drives for personal computers appeared on the market, in about 1983, they were very expensive, costing almost as much again as a floppy disk-based system. It was hard to justify the purchase of such a device for a single user. So a system of cabling and commands was developed to allow several users to share a hard disk. Thus, one of the first local area networks for personal computers, Omninet from Corvus, was created for, and designed around, hard disk sharing.

The need for personal computer LANs was met by adapting software and hardware originally developed for larger computers. Systems such as ARCnet and EtherNet were introduced for personal computers. The hard disk cost justification argument also was applied to letter-quality printers (in 1984 a NEC daisy wheel printer cost almost as much as an IBM PC XT, and a laser printer cost considerably more). Because

a single user was unlikely to keep a printer busy all the time, it made sense to share it among several users on a LAN.

There still are plenty of peripherals—such as color scanners, optical disk drives, tape backup units, and typesetters—that continue to justify LANs. Even though there has been a steady decline in the price per megabyte of hard disk storage, sharing sophisticated drive systems, such as hot-swappable drive arrays, still makes economic sense. Furthermore, an even stronger argument in favor of networks has arisen, based on the potential productivity gains made possible by sharing data across a network.

As personal computers spread through offices, people found that several users would be working on the same data. This led to the "Sneaker Net," which involved running between computers carrying floppy disks of vital data from one user to another. Cabling together the computers into a LAN eliminates Sneaker Net and allows much more sophisticated data sharing. You can see an example of a LAN diagrammed in Figure 1.13. In most LANs, including those based on Novell NetWare, one particularly powerful computer provides the bulk of the storage capacity on the system and controls the network software. This is referred to as the *file server*. Because of the pivotal role of the file server, extra attention must be paid to its security.

Any personal computer working as part of a LAN can be referred to as a network node or a workstation. Note that there is a big difference between a terminal on a multiuser computer and a network workstation. The central computer in a multiuser computer system does all of the computing, but all of the workstations on a network




Figure 1.13 Diagram of a simple LAN.

possess and contribute their own processing capability. While the network file server might appear to be a central computer, it merely serves up data storage and other services to the network workstations.

From a security point of view, the workstation can be both a target in its own right or an entrance to the network as a whole. Consequently, workstations require particular attention when securing personal computers. This is true even on networks that do not have a file server. Such networks, known as *peer-to-peer*, allow each workstation to choose which resources it makes available to all of the other workstations. Examples include Personal NetWare and Windows for Workgroups. You can see the latter at work in Figure 1.14.

WANs

Of course, there are exceptions to many of the definitions presented here. For example, personal computers can be attached to multiuser minicomputers, not just as dumb terminals, but also as intelligent workstations. In such cases, the computing tasks are shared between systems. This often is the case with Unix-based systems where workstations that are powerful in their own right also are connected together to share resources.

Alternatively, a personal computer can itself support several users working at terminals, making it a micromini. Personal computers based on the 80486 and Pentium

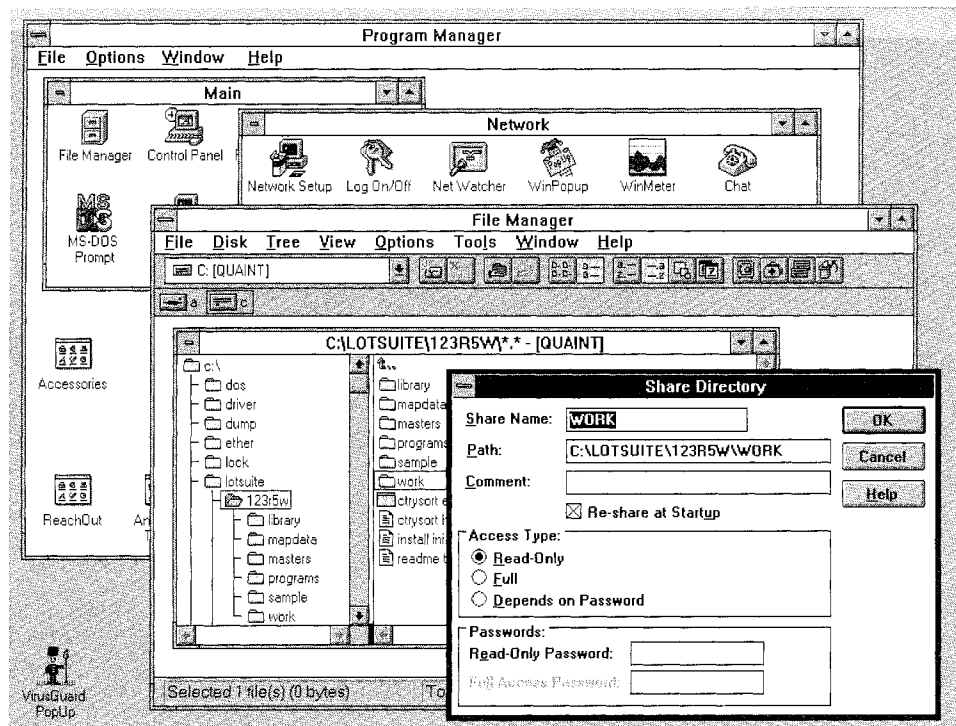


Figure 1.14 Windows for Workgroups provides point-and-click sharing of files between machines

chips can be used in a variety of ways: as incredibly powerful standalone computers, as file servers for large networks, or even as multiuser computers that provide mini-computer level resources to a number of terminals. Indeed, Pentium-based PCs have powers equal to systems that, in 1985, would have taken up a large filing cabinet.

Categories of Personal Computers

This book deals with personal computers in general, rather than just one brand or type. Security is a concern to all personal computer users, whatever model they use. To avoid confusion, the term personal computer will be used in all situations where a general statement is being made. When comments refer to a specific type or make of personal computer, the following terms will be used.

PC

When IBM used the initials PC for its first microcomputer, it created a dilemma. Do we use the acronym PC to refer to all personal computers or just the models that IBM sells? In this book, the term PC refers to any computer, of any brand, that is based on the Intel i86 chip series (8088,80286,80386,80486,and Pentium) plus the various work-alikes from companies such as AMD, IBM, and Cyrix. A parallel definition of PC is "designed to run MS-DOS or Windows." This is the operating system developed by Microsoft, although there now are separate versions from IBM (PC-DOS) and Novell (Personal NetWare or DR-DOS 7).

Mac

The Apple Macintosh personal computer, conveniently and affectionately known as "The Mac," is quite different from the PC in many ways. However, at heart, a Mac is a personal computer. Indeed some Macintosh fans have described it as "the only truly personal computer." While such people might object to the term PC being reserved for IBM-compatible machines, they still will find plenty of Mac-related tips and techniques in this book. Whenever appropriate, attention is drawn to the way specific security issues affect the Mac.



Which Mac?

The Macintosh computer will be referred to as Mac. The term MAC, all in upper-case, will be used for message authentication coding, which is a system of security required by some government agencies.

RISC machines

A new type of personal computer, which can be referred to as an RISC machine, is breaking down traditional computer categories. The acronym RISC stands for reduced instruction set computer, which is a design of processor chip that increases processing speed by reducing the number of commands that the chip understands. Programs written for RISC chips actually have to build up more complex commands

out of the reduced instruction set, resulting in bulkier programs and more commands to be processed, but the RISC processor handles them faster than a CISC chip, or complex instruction set computer. Examples of RISC chips are the DEC Alpha, the MIPS 4000 series, and the Apple/IBM/Motorola PowerPC chip. The hugely successful Intel i86 chip series is an example of CISC technology.

Because RISC technology shifts many basic functions from hardware to software, it facilitates the transfer, or porting, of an operating system from one machine to another. Thus we now have Windows NT, which started out as an advanced operating system for Intel i86-based systems, ported to the DEC Alpha and MIPS chips, with a PowerPC version quite possibly in the works. The most popular RISC systems so far have been the Apple Macintosh PowerPC machines, 1.7 million of which were shipped in the first 12 months of production.



Stacking the DEC

A prime example of what RISC technology is doing to traditional computer categories is DEC's 166 MHz AlphaStation. This comes in a plain desktop box, has standard PCI and EISA expansion slots, built-in Ethernet and audio, yet performs better than some desk-sized minicomputers. A file server configuration of the same system, the AlphaServer 400 41166, is recommended by DEC as a migration path for users of its MicroVax 3100 system. The desktop box can run the same OpenVMS applications as the MicroVax, along with all Windows NT and OSF (Open Software Foundation) server applications.

OS/2

While the terms PC, Mac, and RISC are hardware-based categories, it also is possible to define machines according to the operating system that they use. The first popular operating system for microcomputers was CP/M, which eventually was overtaken by DOS. In 1984, IBM introduced the PC AT, which used the Intel 80286 chip. This chip had computing power way beyond the 8088 chip that powered the original PC and for which the original DOS had been written. While the PC AT and subsequent clones could run DOS and run it much faster than the basic PC, it was clear that a new and improved operating system would be needed to tap the full potential of these machines.

In conjunction with Microsoft, IBM went about developing such an operating system and called it OS/2. Although Microsoft later dropped out of the project, IBM has persevered with OS/2. The latest version is called Warp. It delivers true multitasking, which is the ability to perform more than one task at once (for example, you can be downloading data from the company mainframe while sorting a large database, editing a word processing document, and printing a complex graphic). You can see OS/2 Warp at work in Figure 1.15.

Multiple windows

The Apple Macintosh, which first appeared early in 1984, was the first popular mass-market personal computer with a graphical user interface (GUI, pronounced



Figure 1.15 IBM's OS/2 Warp is a powerful and attractive operating system

goeey). However, there were precedents in the Xerox Star system and Apple's own Lisa. Although Microsoft was making a lot of money in the early 1980s selling DOS, a character-based, command-line interface, it could see the advantages of a graphical approach to user interface design. In 1983, Microsoft announced Windows, and by Fall Comdex 1983, an impressive list of companies had signed on to the "We Do Windows" campaign. It was 1986 before the first version of Windows started shipping.

Most people agree that the first really good version of Windows was 3.1, which didn't arrive until 1990. Even at this stage, Windows was not an operating system in its own right, merely an interface or operating environment added to DOS. With Windows NT, Microsoft finally delivered a complete operating system, widely praised for its multitasking abilities plus its security and administration features when used as a network server. However, NT places heavy demands on hardware and is overkill for the ordinary desktop user.

It is for these users that Microsoft has been developing the product known first as Chicago, then as Windows 95, which is a self-contained operating system that can be installed and used without DOS. In the meantime, Microsoft has promoted Windows for Workgroups 3.11 as the "standard" version of Windows. This installs on top of DOS; however, if you have network cards and cabling, it offers peer-to-peer networking for file and printer sharing.

Unix

Ever since the Intel 80386 appeared in 1986, you have been able to run versions of the Unix operating system on what essentially is a personal computer. A true multi-tasking, multiuser operating system originally developed by Bell Laboratories, Unix is not strictly within the purview of this book; however, references are made when appropriate. For example, some companies are looking to Unix as an alternative to Windows NT and OS/2 and as a way to avoid the problems of local area networks based on MS-DOS.

Unix is used extensively in universities and research organizations and in scientific applications. Unix has the capability of being a very secure operating system, although not all of the security features are used all of the time by all users. It was a Unix-based network, the U.S. DARPA Internet, that suffered one of the most widely reported "attacks" in computing history: the Robert Morris Worm.



As the Worm Turns

Five years ago, I opened this chapter with the 1989 indictment of 24-year-old Robert Morris, Jr. by a United States federal grand jury. He was charged with violating the Computer Security Act of 1987 (the penalty for this crime is up to five years in prison and a fine of up to \$250,000). Morris wrote a "worm" program that exploited soft spots in the Unix operating system (including a bug in sendmail), then released it on the Internet.

Within hours, almost 6000 computers ground to a halt as this 4000-line program replicated itself. The problem could have been much worse, but a hastily assembled team of computer experts identified the problem and instituted countermeasures. Thus was born CERT, the Computer Emergency Response Team, which now monitors the Internet and issues security alerts whenever problems are detected.

Back then the Internet was relatively unknown outside of academic and government circles. The media was largely ignorant of the different categories of malicious software. The Morris program often was called a virus, which sounds pretty nasty, but it actually was a worm, which doesn't sound as exciting (while both are self-replicating, a virus requires a host program, while a worm does not).

Others

Security also is a concern to users of microcomputers such as the Commodore Amiga and the Atari ST. While these systems are seldom specifically addressed in this text, it is acknowledged that they frequently are employed in roles that are important to their users. Much of what is said about risk analysis, protection techniques, and security policies does apply as much to these and other systems as it does to Macs and PCs. In some areas, such as viruses, specific concerns for users of these systems are discussed.

User Categories

Clearly, personal computers have a tremendous range of applications and a correspondingly diverse range of people who can be called personal computer users. While all responsible users will want to protect their systems, there is considerable disparity among users, both in the level of threat and the value of potential damage. The following three categories represent one way of grouping users with respect to their security needs.

SOHO users

These are not denizens of the seedier part of London's West End but the so-called small office/home office users. Identified in the early 1990s as a potentially lucrative market, these users have been buying large numbers of computer systems for such tasks as bookkeeping, education, playing games, or managing one- or two-person enterprises. Most SOHO users work with a personal computer for their personal enjoyment or benefit. Work that they perform with the personal computer is for themselves rather than for their employer. For the most, these users own the personal computers that they use.

Group users

I have used the term *group users* to refer to people who use personal computers as part of their work for an organization. In most cases, these users work on their employer's premises, using personal computers owned by the organization. By using laptops, notebooks, and personal digital assistants, these users might work outside of the office, but such use essentially is an extension of the office.

Supporters

The term *supporters* does not refer to sports fans. It includes all of the people within an organization who support, assist, or manage group users. This includes anyone who has some level of responsibility for the personal computer resources of an organization, from the help desk to network administration, from *ad hoc* experts to MIS staff.

Obviously these categories are not rigid or exclusive. You might well belong to more than one of them; however, in each, you will have a different set of concerns when it comes to security. In writing this book, every effort has been made to address the security needs of all groups, from the private user concerned about viruses on game disks, to the group user anxious not to imperil data belonging to his or her employer, to supporters coordinating the protection of an organization's resources.

Naturally, all of the advice given does not apply equally to all groups. The group users and supporters most often are addressed because they have more problems with which to deal. Some remarks will apply only to the supporters because they have a wider responsibility for security than the other two types of users.

What Is at Stake

Assessing the value of a personal computer system, the tasks it performs, and the information it handles to carry out those tasks is a complex undertaking. Chapter 3 takes a more detailed look at putting a price on information systems and evaluating the consequences of their failure to perform as usual. For now, it is sufficient to say that there are four main aspects of value in any personal computer system:

- The hardware itself
- The software that the hardware uses to process information
- The information being processed
- The system's ability to continue processing

The hardware itself

At first glance, most hardware seems easy enough to value. You probably know, or can find out, what you or your company paid for it. You can assign a value to hardware using normal business concepts like purchase cost, depreciated value, and so on. However, there are some factors that are less obvious, such as replacement value in the case of hardware that is no longer readily available. For example, what if your organization has come to rely on a system based on hardware that is no longer in production? If that hardware fails or is stolen, the process of converting to newer hardware might cost a lot more than the purchase price of the hardware.



Chip Chop

In 1975, some 17% of all chips, based on dollar volume, were used by the military. By 1994, that number had declined to 1.1%. Many of the chips currently in use in military equipment, which often is expected to have a life span of 20 years or more, are no longer in production. In some cases, the companies that made them no longer exist. There now is a whole segment of the computer industry devoted to maintaining stocks of "mil-spec" chips and, where possible, converting commercial chips to military specifications.

The software used by the hardware

The value of software is quite a different story. What you are paying for when you buy software, in addition to the raw material of disks and manuals, is the right to use the intellectual property that they embody. In some ways, this is a defense against theft. Someone could steal your copy of the MegaSpread manual and even your original MegaSpread program disks; however, under typical licensing agreements, you still would own the right to use MegaSpread (for this to have real balance sheet value, you need to be able to prove that at some point you really did purchase the program).



Software Invaluable

According to a report in the *Boston Globe* in September, 1993, about \$10,000 worth of computer equipment and software was stolen from Optimum Manufacturing, Inc. A company spokesman said that the software was much more valuable than the computers, but declined further comment on the weekend burglary. Optimum parts have been used in satellites, tanks, fighter planes, and missiles.

In the case of software that was custom written or internally developed, the question of value is murkier still. A program that is very industry- or company-specific might have little value to other users. If you lose all copies of the program and the developer is no longer in business, the program could be said to be priceless.

The information being processed

While the term *information processing* brings to mind piles of accounts receivable, customer orders, mailing lists, and so forth, the reality for some personal computers is quite different. The processed information includes budget projections, letters, memos, proposals, resumes, and so forth. There are several aspects to the value of whatever information is handled by a personal computer:

- Value to you
- Value to others
- Negative value
- Value of immediate access

The value to you. Suppose that you lose the latest cost estimate worksheet for a major competitive bid. Recreating the worksheet will cost you time and effort. You also might experience a loss in terms of credibility and goodwill within your organization if the missing file means that you disrupt the schedule for submitting the bid. This demonstrates the very personal value of the information to you. If the lost file results in a lost bid, then the wider value of the data to you and your organization becomes clear.

The value to others. Suppose that the file that contains the cost estimates disappears from your personal computer system into the hands of your competitors, who use it to win the bid. This probably is the most obvious demonstration of the value of your information to others. Such data as customer lists and marketing plans fall into the same category. Yet direct competitors are not the only people who might covet the information on your personal computer system. If you are in the business of selling information itself, then there are likely to be those capable of thinking about how to get your information for free.

The negative value. Information that is not of direct value to you or your competitors still might have negative value. Many of us have at one time or another used a personal computer to prepare a nice-looking resumé. If the resumé is intended for a prospective employer, discovery of the resumé file on your current employer's personal computer can be embarrassing to say the least.

Negative value was clearly demonstrated by the infamous Willard Scott/Bryant Gumbel affair. An internal memo in which NBC television host Bryant Gumbel expressed personal observations about co-workers was made public, and much was made of Gumbel's negative assessment of co-worker Willard Scott. The memo actually was written by Gumbel at the request of his boss, but this did not lessen NBC's embarrassment. The memo came to light because it was stored on an insecure computer system!

Other examples of information with potential negative value include internal findings about product safety, employee evaluations, environmental test results, and so forth. Indeed, most internal documents that reflect negatively on an organization or individual have potential negative value.

The value of immediate access. Suppose that you come to work in the morning and your personal computer system does not work the way that you expect it to. The file that you need is not where you thought you left it. In a situation like this, you learn the value of immediate access. The file might not be permanently lost, but access to it is delayed, wasting time and effort, causing a dent in productivity.

The system's ability to continue doing the job

Closely akin to the value of immediate access is the ability of your computer system to keep doing the job. As personal computers increase in processing capability, they are assigned increasingly important tasks. Using personal computers for such tasks as order processing, customer reservations, stock management, and data acquisition means that their role is critical. The cost of system disruption, in terms of lost business and goodwill, can be considerable.

The tasks you perform with your personal computer might not be critical to an organization's profit and loss, but they still might be very important to you. While personal computer security clearly is about letting personal computers get on with what they do, precise statements about the value of personal computer systems are difficult because any respectable list of the different things that personal computers do would be far too long for this book. The approach taken here is to assume that the system with which you are concerned performs tasks, the importance of which is clear to you.

Attacks, Threats, and Scares

To discuss information security effectively, an agreed vocabulary is needed. In addition to the regular terms and phrases of computer technology, the following terms will be used:

attack — General term for any action or event that threatens to interfere with the proper functioning of a computer system or that seeks to achieve unauthorized spread of information entrusted to a computer.

active attack — Action initiated by a person that threatens to interfere with the proper functioning of a computer system or that causes unauthorized spread of information entrusted to a computer. Examples include intentional erasing of files, unauthorized copying of data, or the introduction of a virus designed to disrupt the computer's operation.

passive attack — An attempt to gain information or resources from a personal computer system without interfering with its operation — such as electronic eavesdropping, Van Eck phreaking, or placing a tap on a network, all of which can yield important information about the system as well as appropriate the data that is in the system.

social engineering attack — General term for deceptive practices that attempt to obtain information from people using social, business, or technical discourse. For example, calling a network administrator and misrepresenting yourself to get someone else's password.

incident — When an attack takes place or a threat materializes, you have an incident. Some examples are a power failure or an attempt to delete a protected file.

breach — A successful thwarting of security measures, such as the theft of a PC or the deletion of valuable data files.

threat — Anything that has the potential to interfere with the proper functioning of a system or cause the loss or unauthorized spread of information entrusted to a system. Examples are power failures, viruses, hackers, or careless users.

Do You Really Need This Book?

The purpose of this book is to give you a clear perspective on the problems of security as they relate to personal computers and networks. If you have responsibility for protecting data, either your own or that of an employer or client, this book will help you meet that responsibility with confidence.

A simple test

Have you stored valuable information on a personal computer? Are you sure that it will be there in the morning? Are you sure that it will not appear in the next edition of *The Times* or *The Wall Street Journal*? What *will* you do if your hard disk crashes right now or the office is gutted by fire overnight? Are you happy letting the clerk from the temporary agency update your payroll files or finish typing that competitive quote? Could your competitors use the information that is on your personal computer to their advantage? Are you sure nobody is reaping a windfall from the airwaves in your office, even as you read this?

If you are not comfortable with your responses to the previous questions, then this book probably will help you. Just buying it and putting it on the office bookshelf might help you feel slightly more at ease. Reading it and taking its suggestions to

heart definitely will do more than that. While it might be some time before you can feel confident with the current state of play in the security game, you will know what the threats are and how to defend against them. If you use a personal computer in your work or at home, try taking this simple test to help you to grasp the security implications of computer usage:

1. Do you use password protection on any of your data files?
2. Do you have a surge protector or some other power line conditioner fitted between your computer and the main power supply?
3. Are your unencrypted data files free from any references to personal or company bank account numbers, credit card details, or phone access codes?
4. Could the contents of all of the data files that you created in the last six months be leaked to your co-workers without causing you any embarrassment?
5. Is your computer attached to your desk by anything besides cables and gravity?
6. Do you have a written record of your computer's serial number?
7. Are your fresh data files backed up every working day?
8. Have you chosen a network access password that has nothing to do with the following: your name, the company's name, the names of your children and pets, your date of birth, phone number, or social security number?
9. Have you ever refused to make a copy of a piece of software because to do so would be a violation of the licensing agreement?
10. If someone gave you a free copy of a great new game designed to run on your particular brand of personal computer, would you know how to tell if the disk contained a virus?

Now look at the results. What about number 3? Are you unsure about what unencrypted means? It means unprotected by a password so that just about anyone who comes across the file can read its contents. (I remember reading a simple word processing file that listed company credit card numbers on a floppy disk that was being thrown away!)

If you answered "No" to three or more of the questions, you have a lot in common with the typical personal computer user. Your security awareness needs improving, and you could do a lot more to protect yourself and your data. This book will tell you how.

If you answered "Yes" to six or more questions, then you are either very security conscious or naturally suspicious, but you still should read this book. You will want to be sure that you have considered all the angles, not just those in the previous questions. Unfortunately, generating a comprehensive security plan often is an exercise in applied suspicion.

Only by taking a dim view of people can you be sure that your personal computer and the data entrusted to it are free from risk. Hopefully, you will take such a view temporarily and not become completely jaundiced. One way to increase security awareness for yourself and others within your organization is to use a program like SAFEware, which is shown in Figure 1.16.

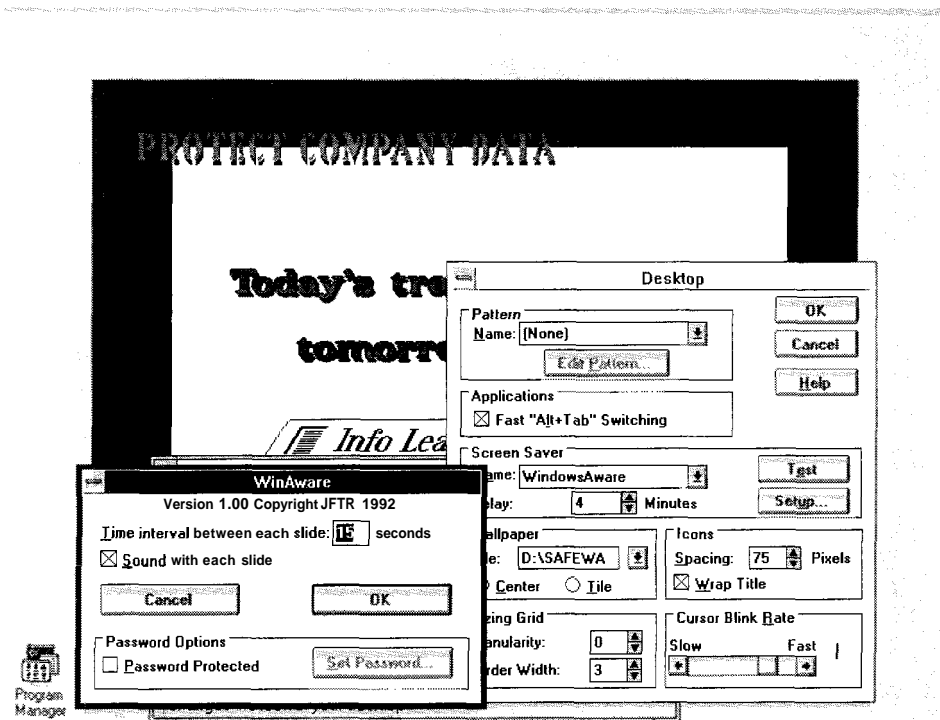


Figure 1.16 Your computer can be used to increase computer security awareness, using a program such as SAFEware

A tougher test

The next test is harder. Read the questions before putting pen to paper. The aim of this test is to help you put matters into perspective rather than elicit accurate answers:

1. Without looking at the computer, list all of the programs that are on the computer's hard disk (if you do not have a hard disk, then try the nearest 20 floppy disks to the computer).
2. Add to the list descriptions of the data files and documents that are on your hard disk or in your collection of floppies (you do not have to list exact file names, but enough of the name and/or contents to identify the file).
3. Write down the date and time of the last complete backup of your hard disk (if you do not have a hard disk, write down the last time that you made backup copies of your six most important data disks).
4. Name three popular programs for your computer that feature password protection.
5. At 3 A.M., how many locked doors are there between your computer and the street?
6. Does anyone in your organization who uses a computer feel disgruntled, disaffected, badly done by?

7. Are you sure that you know what is in all of the data files that are on all of your disks?
8. If you are using a network and call your network administrator to say you have forgotten your login password, will he or she issue you a new one over the phone?
9. Do you know how many modems are on your network and how many of these are set to auto-answer incoming calls?
10. If I walked into your office, would I be able to find the keys to keyboard locks sitting in unattended machines?

By now, you should see the point of these tests is to increase your awareness of the size of the problem without causing undue alarm. The fact is that very few people:

Know exactly what is stored on their computer.

Perform backups in a timely fashion.

Use password protection correctly.

Make proper use of all of the locks that exist between their data and the outside world.

A frightening possibility

To further assess where you stand in relation to personal computers and security, check how you relate to this scenario: Tomorrow morning you arrive at work to find the office in chaos. There has been a break-in. Your PC has been stolen. What is your first thought?

Like many people, you probably will be racking your brains trying to remember what exactly was stored on the computer. While some of us are responsible enough to keep records of our equipment serial numbers, few of us have a clear and up-to-the-minute recollection of what is stored on our computers at any given time. Damage assessment is difficult when you are not sure what exactly is missing.

The next thought probably is "When was the last time that I made backup copies of my work?" This thought is quickly followed by "Did those get taken as well?" If, like many people, you made your backup onto floppy disks that sit in a disk box next to the computer, the answer might well be "yes."

The next step is to consider how to get on with your job, given that disk files of documents, statistics, ideas, proposals, contracts, estimates, invoices, accounts, and even amorous jottings now are "out there" rather than safely tucked away on your hard disk.

These days, a computer rental is just a phone call away in most cities. When the rental unit arrives, it is time to start to think about backup copies of your favorite programs. If these were not stolen, you still might have to go through lengthy installation procedures to get them up and running. You might have had a lot of small utility programs on your PC, batch files and so forth, that made life easier. Do you have copies of them?

Even when you are up and running, you still are not out of the woods. Did they steal the printer? Does the software that you are using work with the replacement printer that you rented? If they left the printer, **will** it work with the new computer?

Did they leave that custom printer cable that lets that Compaq 486 print on your Sukidata 180 printer (the cable that Fred what's-his-name soldered up for you before he left to be a computer consultant in Tasmania)?

After several days of incredible frustration, you might get a call from the local police. They have recovered your PC! They cannot tell you whether anyone has used it or not, but you're glad to get it back, right? Maybe. You already have canceled all of the credit cards, whose account numbers, expiration dates, and current balances were stored in that personal money manager program that you got for Christmas. Your bank has been warned to watch out for checks on your account, details of which were in that same handy program.

You have failed to meet the deadline for bids on that lucrative new contract because the cost estimates were on the computer, not on safely stored floppies. If anyone has read that innocently named word processing document called MYNOTES, your private life could be in for stormy weather. As you drive across town to pick up your recovered PC, you realize that if the culprits have damaged your hard disk, what you are going to collect is a fairly worthless piece of iron.

Lessons learned

This scenario would sound a lot like scare tactics if it was not repeated every day in so many places around the world. Fortunately, it is only a scenario, and several lessons can be learned, all of which **will** be reviewed in greater detail elsewhere in the book:

- Make the office more burglar-resistant.
- Fix your PC to the desk.
- Do the same for printers, cables, and other accessories.
- Make frequent backups of data and programs.
- Keep the backups in a safe place.
- If there is a system startup password facility, use it.
- Use the password protection on programs that offer it.
- Use file encryption software on all other sensitive files.

The fact that some of these lessons are fairly obvious does not make them any the less valuable. Security requires common sense more than expensive equipment and commitment more than complexity. For example, your company might have a policy about storing personal data on company computers or locking away backup disks. However, unless it is relatively easy for employees to comply with such policies and someone is charged with enforcing them, they are as useless as a slice of floppy disk.



Disk Asked Her

One infamous instance of sliced disk also is a good example of information having "negative value." An Air Force officer used his word processor to write a letter to his lover asking her to hire a hit man to kill his wife. While military police were interviewing the officer at his desk, he actually sliced up the floppy disk that contained the letter using a pair of pinky shears. Thanks to some ingenious reconstructive electromagnetic surgery, the disk, and its vital evidence, was salvaged and formed the basis for an arrest and conviction.

Further Questions of Security

This chapter has asked you to consider how much is at stake when you employ personal computer technology in your endeavors. The fact that you have picked this book up suggests that you already are concerned about issues of data security. You are aware that there are problems and want to know more about them. You probably want to prevent security issues from eating into the benefits you have gained from personal computers.

Many people are alarmed by the kinds of questions that have been raised in this chapter. However, these are the questions that you should be thinking about if you are in any way responsible for information handled by personal computers. Fortunately, this book is an antidote to alarm. When you have read it, you still will be concerned, and you might realize that you have a lot of work to do to make your data secure, but you will be armed with a realistic understanding of the threat to your data. You will know how to analyze the risks and dangers and how to formulate an appropriate response. You will be able to assess new threats as they arise, because you **will** have a clear, security-conscious picture of personal computers, and how they handle information.

Questions of security have graduated from the periphery of the personal computer community to the board rooms and management meetings of every company that uses personal computers. They now represent serious anxieties, the sort which, if left unresolved, threaten to discourage prospective users who stand to gain so much from the power of microcomputers. Three basic questions serve to place the subject matter of this book in perspective:

- What do you stand to lose?
- What are the sources of danger?
- How can you protect the former from the latter?

What you stand to lose

Several ways of looking at this question were presented in the earlier section, but the question needs further consideration. In 1987, the U.S. journal *Government Computer News* asked the NCSC (National Computer Security Center) this question: "What are the greatest security problems facing typical government agency informa-

tion systems managers?" The reply was unequivocal: "The lack of awareness among computer users." According to a survey conducted in 1994 by *Infosecurity News*, things have not changed a whole lot in seven years. Readers of the magazine cited the lack of end-user awareness as the single most important security concern. Other leading obstacles to achieving adequate infosecurity levels are charted in Figure 1.17.

A major reason for this lack of awareness is the failure to grasp what can be lost through security breaches. The good news is that awareness is a lot cheaper to raise than the cash required to install additional security devices that will be ineffective if users still are complacent. By stressing what can happen to profits and productivity because of lax security, employees can be motivated to attain higher levels of awareness. Indeed, it is not unreasonable to make the connection between job security and information security.

Experience has shown that computers that handle accounts stand to lose money. Computers that handle valuable data stand to lose data or have it devalued by unauthorized distribution. Computers entrusted with sensitive information have the potential to cause tremendous embarrassment if their security is compromised. When it comes to information that is vital to society, we stand to lose a lot from what the NCSC calls the "proliferation of untrusted technology." Lax personal computer security in areas such as health and safety puts more than mere data at risk.



A Child in Chains

When a clerk at University Medical Center, Jacksonville, went into work last Sunday, she took along her 13-year-old daughter, Tammy. Taking advantage of poor computer security, Tammy obtained a two-page report listing former emergency room patients and their phone numbers. She then proceeded to call people on the list and tell them that they had tested HIV positive. Appearing in court this week in handcuffs and leg shackles, Tammy was ordered into state custody. The judge justified harsh measures in this case because Tammy "seemed unconcerned about her arrest or the possible effects of her actions."

Orlando Sentinel, March, 1995

The sources of danger

This book should answer the many questions that personal computer users have about the safety of the data that they have entrusted to a machine that they now realize is fallible and open to attack. For example, does the new telephone connection for your personal computer increase the chance of your files being infected by a computer virus? How likely is it that your personal computer network **will** be attacked by a hacker? Is it true that data being displayed on your computer screen can be picked up by equipment outside your office? Does the government really try to prevent companies from using any secret codes that it cannot break?

Questions like this have received scant attention during the headlong rush to harness the potential benefits of the microcomputers. Those who raised such questions risked being classified as paranoid, reactionary, or possessed of overactive imagina-




Figure 1.17 Chart of perceived obstacles to information security (Used with permission of Infosecurity News)

tions. Nowadays, thanks to increasingly well-publicized and ambitious feats of computer tampering, these questions are being asked among rank-and-file PC users.



A Place Where We Can Meet

The NCSA Information Security Forum on CompuServe provides a convenient place for serious, moderated discussions with fellow security professionals. This can take the form of messages posted to the dozen or so bulletin boards, which are broken down by subject. You also can chat informally with other users who are online at the same time (see Figure 1.18). In addition, special conferences are held with guest speakers who can answer questions online.

Unfortunately, the questions play better than the answers. Tales of teenage whiz-kids penetrating multinational data networks get more media attention than the simple list of administrative procedures that could have prevented the incident in the first place. When an outbreak of a computer virus makes national news, people tend to see viruses in every computer malfunction and error.

The problems of viruses and attacks by hackers have overshadowed most other aspects of computer security. While researching this book, it was possible to find a lot of headlines like this: "Defense Department Computer Broken Into Again" or "Virus Threat Set to Worsen." It was harder to find anything along the lines of "Thieves Guess Password, Divert Millions" or "Stolen PC Contained Company Secrets."

Companies do not like to report security breaches that could easily have been prevented, and the fears of computer users are fed by media hype, overshadowing the

mundane subject of preventative measures. In some ways, this situation is natural. There is glamour and excitement (of sorts) in talk of computer hacking, vandalism, sabotage, fraud, and theft. The other side of these issues is the less than glamorous talk of procedures, rules, audit trails, and paperwork. As I'll demonstrate in the next chapter, you can greatly increase the security of the data on your personal computer with a few simple procedures that cost n o t h g to implement.

As we come to rely upon personal computers to carry out an ever-increasing share of the information management tasks w i t h our society, you should realize that the underlying design of these systems has evolved little from the early days. The first personal computers were spawned by a desire to open up computing to the masses, not keep people away from sensitive information.

By their very nature, personal computers are "open" to all users. With few exceptions, today's typical personal computer can be turned on by anyone, operated by anyone, opened up by anyone, and carried off by anyone. The number of people who know how to operate a computer is growing all of the time, as is the number of people who can write programs for them. Clearly the personal computer is made to be used, designed to be accessed freely. Against this must be balanced the growing need to restrict access.

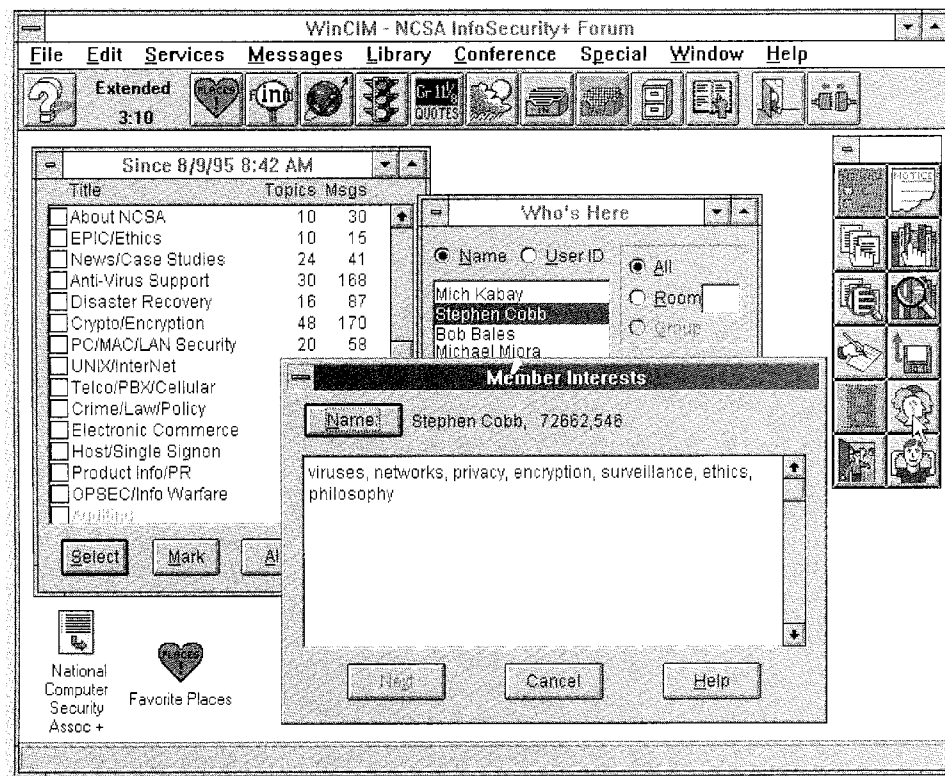


Figure 1.18 While checking messages on the NCSA Information Security Forum, you can look to see who else is online for live conversation.

When it comes to determining the possible threats to your system, your best tool is a healthy imagination. Ask yourself what could possibly go wrong and in what ways could you be ripped off. When it comes to determining which of the possible threats are probable, you need a healthy dose of realism. Make an objective assessment of the question: "Is it worth the risk for someone to try this?" To stir your imagination, consider these scenarios:

Scenario 1. Your organization establishes a "work at home" program in which employees with personal computers at home use modems to dial up the computers at work, receiving and completing assignments over the phone lines. An employee's wallet is stolen, but she is embarrassed to admit that the phone number and password for logging on to the company computer were in the wallet, along with her business cards. An enterprising felon now has a valuable item to peddle, and a computerized felon has access to the company computers. The results can mean damaged, erased, or stolen data sold to the highest bidder.

Scenario 2. In an absent-minded moment, an executive with BIG Corp. leaves his unlocked briefcase on the commuter train. An unscrupulous fellow traveler takes the briefcase and the laptop personal computer that it contains. Reading the unprotected word processing file describing plans to take over SMALL Corp., the unscrupulous fellow has several options ranging from seeking a "reward" for return of the computer, to blackmail of BIG Corp., to a valuable tip for SMALL Corp.

Scenario 3. Another executive with BIG Corp. leaves his unlocked briefcase on the commuter bus. An unscrupulous fellow traveler takes the briefcase and the laptop personal computer that it contains. Reading the unprotected checking account spreadsheet, the unscrupulous fellow leaps at the chance to blackmail the executive who has listed several payments to a woman who is not his wife.

Scenario 4. By using computer resources, many types of fraud are made more plausible. In one case, a sales manager used convincing but phony letters on fabricated letterheads to get bogus suppliers placed on the company master ledger. He then authorized "goods received" notices and payment of the phony invoices to the drop box addresses he had set up on the letterhead. By combining a high-resolution image scanner with a laser printer, you can fabricate what appear to be exact photocopies of official documents, forms, and statements.



Risky Business

If you find it hard to look on the dark side and want to give your imagination a boost, check out the Risks Digest. This is a potted version of a moderated Internet discussion group that considers risks arising from computers technology in the very widest sense. Copies of the digest can be found in the library section of the NCSA Information Security Forum on CompuServe, or you can access them via the World Wide Web (as shown in Figure 1.19). You also can subscribe to the digest to receive the latest issues emailed to you on a regular basis. For more about security risks in the wider sense, see chapter 14.

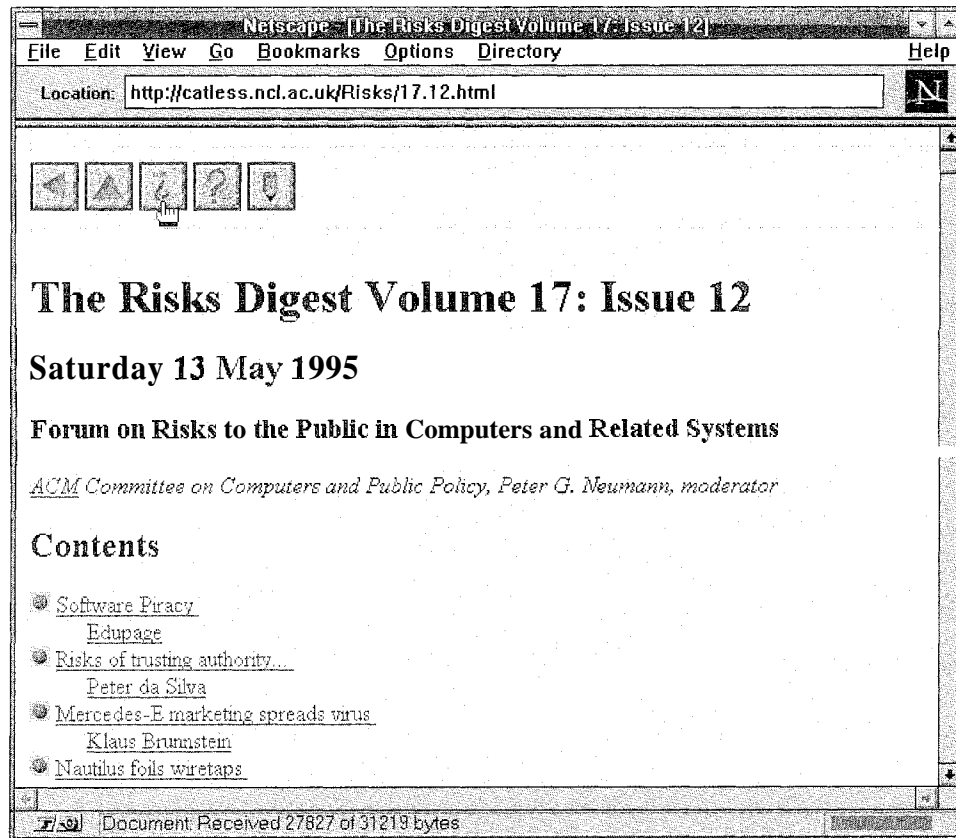


Figure 1.19 You can access the Risks Digest on the Internet via the World Wide Web [http://catless.ncl.ac.uk/risks).

How to protect what you stand to lose

While the real threats are many and varied, nothing is gained by retreating into paranoia. This book is based on the premise that security is best achieved through a balanced reaction to realistically assessed threats accompanied by good disaster-recovery planning.

When security measures go too far, they become burdensome and threaten to choke off the benefits to be gained from using computers in the first place. On the other hand, in situations where security measures are nonexistent or too lax, management is bound to face, sooner rather than later, a rude awakening as to the realities of computer crime and computer-inflicted disruption of regular business. While every possible threat to your systems cannot be foreseen, advance planning will help you make the best possible recovery from an attack.

This book also stresses the concept of an "appropriate response" — finding the right level of precaution to make your data safe, without making your data too difficult to get to. For some users, this simply will mean following a few simple rules.

While no PC user can afford to be complacent, in many cases, there is little need for expensive equipment to safeguard data.

The Network Connection

While there are two chapters in this book that deal exclusively with network security issues, all of the chapters are relevant to networking. This is because most of the computers that constitute today's networks are personal computers. Although some of these personal computers now are extremely powerful and some, like the so-called super servers, are not intended for personal use, the fact that they have their roots in personal computing sets them apart from minicomputers and mainframes. Consequently, many of the security practices and paradigms that were developed to deal with minicomputers and mainframes do not apply or can be applied only with considerable adjustment and/or difficulty.

Securing today's networks of personal computers requires a new approach, one that combines traditional wisdom with fresh insight. We might call this the desktop security perspective. To help network managers and users develop this perspective, each chapter will contain a section called "The network connection" to highlight the ways in which the subject matter of that chapter applies to networks.

Summary

This book attempts to present the problems of data security realistically and to offer practical solutions based on a balanced approach. When suggestions for action are made, they are presented with a full knowledge of what it will take to implement them in a real office, working with real people.

This book assumes that you do not want to erect so many defenses that your data become virtually inaccessible. At the same time, it is assumed that you have data that needs preserving and protecting.

A further assumption is that, after reading this chapter, you realize that there is a problem. You know that information vital to companies, institutions, and individuals has been entrusted to personal computers. You realize that controlling access to this information is just as important as getting it organized and computerized in the first place.

This book does not assume that you are a computer wizard or a power user who can write assembly language subroutines over lunch, but it does assume that you have some idea of the capabilities of personal computers and the roles that they play in enterprises and institutions.

In the next chapter, you can read in detail about how personal computers work, which should help you to protect them.



Whoever does
your wiring
should be
highly reliable.

By Stephen Cobb

Security Solutions

Basic Concepts and Techniques

"Opportunity makes a thief"
FRANCIS BACON, 1598

This chapter aims to raise awareness in two areas: technical and social. I don't think you can feel confident defending data unless you understand the basic workings of the hardware that you are using. This doesn't mean that you have to go out and get a degree in electronics, but you do need to know what happens when you turn on a PC, in terms of BIOS and boot sectors, if you are going to defend against something like a boot-sector virus. The more serious consequences of a lack of knowledge in this area include both over-confidence and blind panic. However, because information is stolen by people, not computers, this chapter also begins to develop the security "mindset" that you will need if you are to make realistic assessments of the threats posed to your personal computer resources.

This chapter also makes some suggestions as to how you can use security resources that you already might have, or which you can acquire at little or no cost, to give you a head start in securing your personal computer facilities. Some of the suggestions are less than elegant, but then again, they won't break your budget. The point is, you can get a lot of security from knowing how personal computers work and applying a good dose of common sense when working with them. Security is as much a question of outlook as it is of outlay. There is no point spending money on security measures if you do not use them, and there is no better place to start than by using the resources that you already have.

People, People, People

You might be surprised to see people at the top of the personal computer security agenda. After all, we are talking high-tech problems here. However, security also is a people problem. Computers are a human invention, and if they don't work properly, it is humans who get hurt—sometimes financially, sometimes even physically.

Machines versus people

In America, the National Rifle Association is fond of pointing out that guns do not murder people; people murder people. While this truism grossly over-simplifies the ethical dilemmas of gun ownership and control, it nevertheless makes an important point about technology in general. For the most part, the impediments to successful computing are human in nature, and not technical. If everybody woke up tomorrow thinking that it was wrong to mess with other people's information, then the next edition of this book would have fewer pages and sell less copies. Until this happens, the sad fact is that effective implementation of personal computer security requires you to take a dim view of human nature.



Criminal Acts

Sometime on Thursday night, a \$4000 Macintosh computer system that 12-year-old Kara Johansen, who cannot speak, uses to do her homework and to communicate with friends, was stolen from her house on King Street in Cohasset.

Boston Globe, January, 1994.

Of course, it is not the purpose of this text to determine whether people are essentially good or evil. However, we can state categorically that human beings are, to varying degrees, devious, grasping, and downright rotten. To secure yourself and your organization against such character traits, you must be prepared to think like your adversary and realize that sometimes he or she is a miserable specimen (there are times when the expert in personal computer security is sorely tempted to agree with Sartre when he said "Hell is other people.")

Risks, resources, and responses

People might be your biggest risk, but they also are your best resource. Successful personal computer security depends upon good people more than anything else. If you have motivated, diligent, and careful people working with you, then a greater degree of security can be obtained and with less hassle. Without the cooperation of the people involved, even the most sophisticated security devices **will** fail. With cooperation, high levels of security can be obtained without recourse to expensive equipment.

What is clear from experience is that both management and employees have security responsibilities. Working together, they can be very effective in protecting the company's computer resources. The British Government's 1993 Audit Commission report on computer abuse noted several areas of management responsibility:

The chief executive and management board must be determined to instill an awareness of the importance of computer security and be prepared to act when breaches occur.

Line management in user departments must ensure that access to facilities complies with the organization's standards.

IT (Information Technology) management must assist in defining cost-effective controls that protect the data that it is holding and processing on behalf of others. It should educate users in the need for controls over computerized data.

The British term *IT management* is roughly equivalent to the American term MIS in this context. The same report identifies one other key element within management: internal auditors. They have a responsibility "to test and advise on the adequacy of security and controls . . . bring to management's attention any shortcomings . . . and emphasize the risks of all forms of computer abuse."

The following are six quick and relatively immediate steps that management can take to improve information security:

Step 1. Post security rules. Write up a basic list of rules that users should follow to preserve security. Display them prominently. Issue a memo to each employee that states the same rules in a form that can be kept handy. Consider posting a copy of the rules on each personal computer. (See Figure 2.1 for an example of what these rules might look like.)

COMPANY COMPUTER SECURITY RULES

ALWAYS:	NEVER:
1. Back up important data files.	1. Use obvious passwords, such as your name.
2. Use your access password.	2. Write down, or be seen entering, passwords.
3. Use screen saver and keyboard lock.	3. Share or reveal passwords.
4. Label all media and lock them up.	4. Use same password for different systems.
5. Check the ID of outsiders using PCs.	5. Boot a PC with a floppy disk in drive A.
6. Report suspicious activity.	6. Install unlicensed software.
7. Scan floppy disks for viruses.	7. Make unauthorized copies of software.
8. Log off unattended workstations.	8. Leave unprotected modems in answer mode.
9. Protect keys and passwords.	9. Leave unlocked computers unattended.
10. Assume someone somewhere is interested in stealing, damaging, or destroying your data and equipment.	10. Assume that nobody is interested in stealing, damaging, or destroying your data and equipment.

Figure 2.1 Sample rules for safe computing

Step 2. Announce computer security enforcement. Let people know that compliance with the security rules **will** be checked and considered part of job performance.

Step 3. Announce computer security incentives. Let employees know that compliance **will** be rewarded. Remember that people who breach security have an incentive to do so, even if it is just boredom.

Step 4. Open a computer security hotline. Let people know where they can get advice on security issues and give anonymous tips about security violations or suspicious activity.

Step 5. Issue computer security alerts. These can highlight specific weaknesses that have come to management's attention and stress the ongoing need for vigilance in maintaining security. In some cases, you might want to issue alerts in response to news items involving security breaches, but beware of giving employees specific details.

Step 6. Appoint a computer security officer. This raises awareness of the issue and sends a clear message about the organization's commitment to security. You don't have to hire a new employee for this position, but you do need to back the person you appoint with the appropriate resources and authority.

These steps are just a starting point. They are not meant as a substitute for an in-depth risk analysis, security policy promulgation, and implementation program (these items are covered in the next chapter). The idea is to make big strides towards raising security awareness without spending a whole lot of money or scheduling a lifetime's worth of meetings. While the previous steps might sound like token gestures, they nevertheless can be very effective. For more detailed discussion about the human factors in computer security, see chapter 14, which suggests ways to improve security through a variety of personnel management tactics.

Backup, Backup, Backup

Backup is not only the best but also one of the cheapest of the available antidotes to any kind of data destruction, whether it is caused by hardware failure, malicious hacking, or simple human error. In its simplest form, a backup is just a copy. For example, by copying programs and data from one disk to another, you double the chances that they will survive an attack. Make a third copy and the odds are further increased in your favor. If you store the backup copies in a safe place, you can get back to work fairly quickly, even if the rest of your computer system is stolen. It goes without saying that you should back up your work regularly and store the backup materials safely.



Free Backup Software

Just about every operating system comes with its own backup and restore software. These used to be pretty rudimentary, but recent versions of DOS have included more comprehensive backup facilities. For example, with MS-DOS, you get both MSBACKUP and Microsoft Backup for Windows. The latter can be seen in Figure 2.2.

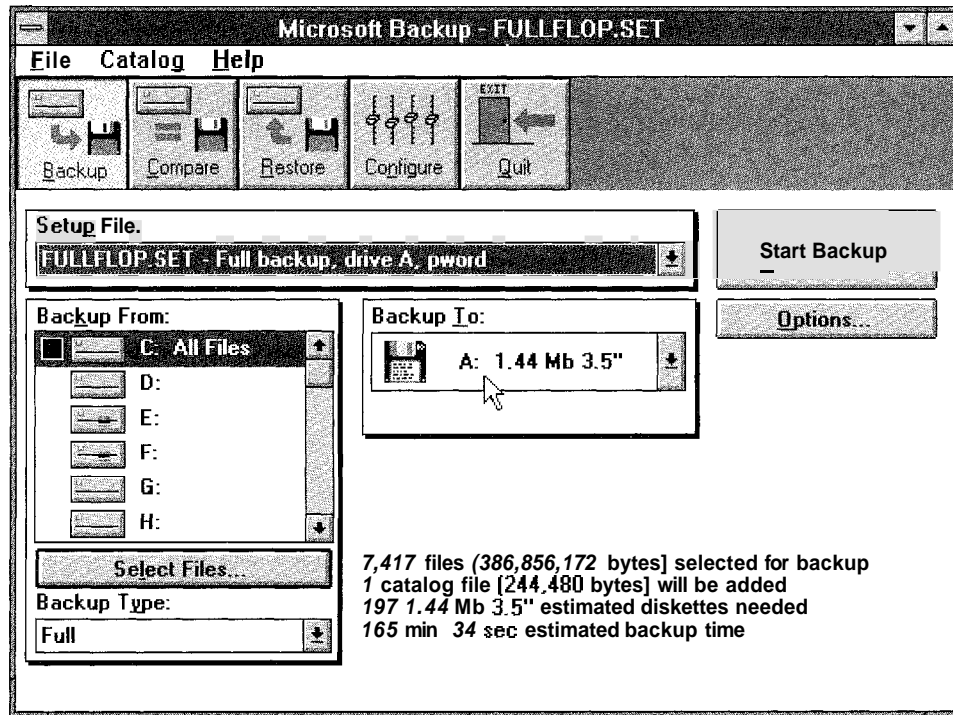


Figure 2.2 Recent versions of DOS have included free backup software, such as Microsoft Backup for Windows

Unfortunately, backing up is a boring task, even if you have a relatively fast tape drive at your disposal. Thus, there is a natural tendency to neglect this task. If you are supporting other users, the single most effective security measure that you can take is to insist on regular backup. If you are a private or group user, the biggest favor that you can do yourself is to do your backup diligently. There are ways to make backup easier and faster, and these are covered in chapter 8, along with discussion of how much to back up, how often, and how to restore from backups when your primary data is lost. Later in this chapter, there will be tips about physically protecting your backup, along with the rest of your system.

Turning on Your Computer

Now that you have an overview of both the major cause of problems (people) and the major antidote (backups), I am going to focus quite narrowly on what might seem like minor details. However, the next few sections, which review the way in which personal computer systems operate, provide valuable information, without which it would be impossible to fully understand the security risks inherent in using personal computers. By getting to know the gaps in the armor, you will be better able to thwart attacks.

Booting up

One of the aspects of personal computing that is both entertaining and frustrating is the use of strange and obscure terms. A typical example is boot disk. Simply put, this is a disk that contains the key components of the computer's operating system, without which the computer cannot do anything useful at all. It gets its name from "booting the computer," which in turn means getting the computer started.

Some frustrated users have assumed that this term came from the desire to kick the darn thing. In fact, boot is short for bootstrap, as in "pull yourself up by your own bootstraps." In a typical personal computer system, the software that controls the hardware, known as the operating system, is provided on a disk. Without software to tell it what to do, the hardware is useless. Somehow the operating system software has to get from a disk into the computer's memory, where it remains for the duration of a computing session.

In fact, a small amount of software can be stored in the hardware itself so that, when the computer is switched on, it has enough sense to look for the disk bearing the operating system. The first part of the boot disk that the computer reads is called the boot sector. This contains the information needed for the computer to correctly read the rest of the operating system into memory.



Reading and Writing

The term read is used to describe the process of retrieving information (for example, reading a file from disk into memory). The term write is used to describe the process of putting information in place (for example, writing a file from memory onto disk). The part of a disk drive mechanism that reads files from the disk and writes them to the disk is known as the *read/write* head.

POST, BIOS, and CMOS

By examining what happens when a computer is turned on, you can learn a lot about how the system can be protected from unauthorized access. When you flip on the switch and power starts flowing into your personal computer system, the first actions are carried out by the BIOS, which is short for basic *input/output* system. The BIOS of a personal computer is the fundamental set of instructions for that particular computer's hardware, including such details as how data is to be presented to the central processing unit, the CPU chip.

The BIOS is stored in something called ROM, an acronym for read-only memory. In this case, ROM is a set of chips on the computer's main circuit board, which is referred to as the motherboard. These ROM chips remember a fixed set of information and allow you to read that information. You cannot write to this ROM. In other words, ROM does not let you change the information that it is remembering. The BIOS in the ROM of your computer performs the memory check and other diagnostics, known as the POST or power on selftest, which happens before the computer tries to read the operating system from disk (see Figure 2.3).




Figure 2.3 Diagram of basic computer architecture ROM/RAM disk

One of the first actions performed by the POST is a check of the keyboard connection. If no keyboard is attached, a warning beep is sounded and an error message is displayed on the screen. Another area checked by the POST is the computer's RAM, or *random access memory*. Because it allows changes to information to be transferred at the speed of electricity, RAM is where programs and data are placed while the computer is computing. Because both RAM and ROM hold information, you measure their size in *bytes*, which is the normal unit of measure when dealing with data.

When you turn on a PC, you often see a series of numbers in the top left of the screen. These represent the size of the computer's RAM as the chips are checked out by the POST. If there are any problems with RAM, the POST instructions in ROM provide you with an error code. There also are error codes for bad keyboard connections and other problems, such as the absence of a boot disk, which is something I will talk about in more detail later.

If the POST reports that everything is okay, then the BIOS is ready to turn over control of the hardware to the operating system. Where does it find the operating system? Normally it is on a disk, so the BIOS looks to a disk drive. On a Macintosh, you can tell if you have reached this point because either a picture of a disk appears with a question mark in it, meaning there is no operating system present, or a smiling computer icon appears to let you know that the operating system has been located.

In most personal computers, ROM actually is a computer chip or collection of chips into which instructions have been permanently etched. Unlike the chips that make up your computer's RAM, which are designed to facilitate changes to the information that they hold, the ROM chips do not need electrical current to store their instructions. When you turn off the power to your computer, the instructions in ROM

remain, while any information in RAM is wiped out. When power is supplied to the ROM chip storing the BIOS, the instructions are carried out. However, the instructions in ROM are limited. On most personal computers, including Macintoshes, they do not include the operating system. In Figure 2.4, you can see a motherboard and the physical location of RAM, ROM, and something called CMOS, which I will discuss in a moment.



Exceptions to ROM Rules

Over the years a number of computer systems have used ROM for more than just the BIOS. The first Hewlett-Packard laptop PC relied on ROM for both DOS and the Lotus 1-2-3 spreadsheet and today's HP Omnibook subnotebook also uses ROM for applications (ROM takes up less space than a disk drive and has no moving parts). Toshiba used ROM for DOS in its T1000 laptop computer. Tandy used ROM to store DOS in desktop machines, together with a small suite of programs. In recent years, some hardware has been fitted with "FLASH ROM," the contents of which actually can be changed using special software. Some modems use FLASH ROM so that they can be upgraded when new communication standards are released.

In addition to the POST, the ROM in PCs stores a program called SETUP, which allows you to define hardware parameters, such as hard disk type, specific to each ma-




Figure 2.4 A motherboard, showing physical location of RAM, ROM, and CMOS.

chine. This is data that your PC needs from session to session but that can be changed (for example, when you add a new hard drive). You can see an example of a SETUP program in Figure 2.5.

Because ROM cannot be changed, the PC uses something called CMOS to store this information. CMOS is a special form of RAM chip that requires very little electricity to retain information (CMOS actually stands for *Complimentary Metal-Oxide Semiconductor* and is powered by a battery). Unfortunately, the battery that powers the CMOS can wear out, resulting in the loss of information and a message such as *Invalid configuration setting, change selection (Y/N)?* when the computer is turned on. Many computer stores sell replacement batteries.

BIOS versus OS

A computer operating system is a program consisting of commands, routines, and conventions that together can be used to run a computer. Essentially an operating system manages the flow of information in and out of the computer, controlling the four phases of operation: input, processing, and storage, and output. While the BIOS sets some of the basic rules for the hardware side of the computer, the operating system provides a complete set of rules and tools with which software can manipulate data.

It is the operating system that controls the opening of files, their arrangement on the disk, the entry of data, the closing of files, sending files to the printer, and house-keeping tasks like copying and deleting files. Because it does so much, the software that forms the operating system is much larger than the BIOS. This is one reason that most personal computer systems store the operating system software on disk and only read sections of it into RAM as they are needed. (Another reason is the ease with which a disk-based operating system can be changed, versus the difficulty of changing an operating system permanently stored in chips.)

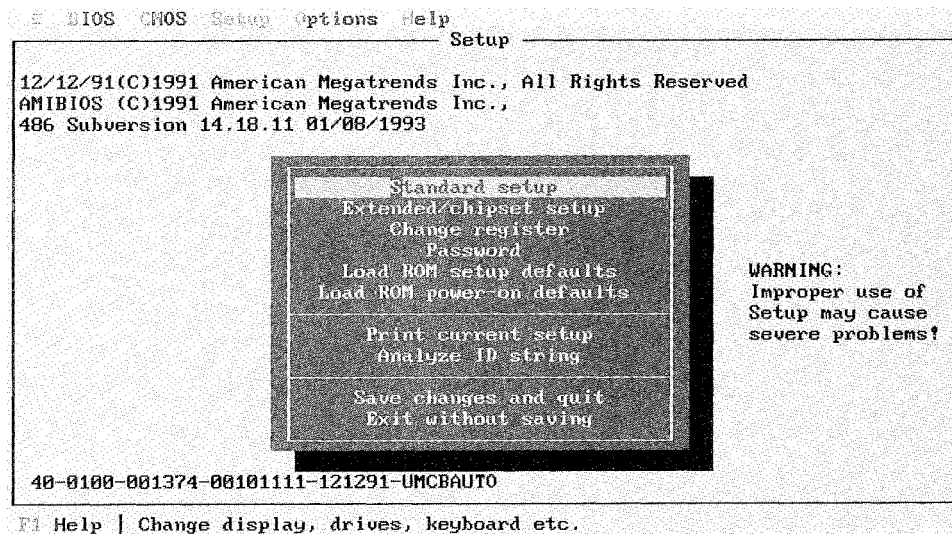


Figure 2.5 The SETUP program, which is used to adjust the BIOS settings stored in CMOS

The fundamental programming that has been done in the operating system makes it much easier for software developers to create applications. Programs like WordPerfect are called applications because they apply the power of the computer to useful tasks. Such applications rely heavily on the operating system to come between them and the raw hardware and BIOS. For example, when you save a file with WordPerfect, it actually is the operating system that carries out the command, deciding where on the disk the file should be placed and keeping track of the location so that it can be retrieved at anytime. You can see this relationship diagrammed in Figure 2.6.

The most widely used personal computer operating system is called DOS, short for *disk operating system*. The word *disk* is used because the operating system resides on a disk rather than in the system's BIOS. Personal computers based on the i86 chips normally use one version or another of DOS, but they don't have to. The same hardware platform can run several different operating systems, simply by booting from different disks (the term *hardware platform* is used to refer to a set of hardware standards). In some respects, a personal computer built on the IBM PC standard is a generic computer, run by whatever suitably prepared operating system you supply on the boot disk (besides DOS, your choices include Unix, Windows NT, and OS/2).



DOS of the Day

These days, DOS comes in many flavors, the most widely known being PC-DOS and MS-DOS. The term PC-DOS actually is IBM's trade name for the version of DOS that they sell. The original creator of DOS is Microsoft, hence the term MS-DOS, which is Microsoft's trade name for its own version of DOS. Microsoft licenses DOS to a number of computer makers, such as Compaq, who get to call their own versions by names like Compaq-DOS. In the late 1980s, a compatible but different DOS was created by Digital Research, the company whose CP/M operating system was rejected by IBM in favor of Microsoft's offering. Known as DR-DOS, this formed the basis of Novell DOS when Novell took over Digital Research.

Drive time

Having seen what happens after you turn on a personal computer, it is time to see how the disk-based operating system comes into play. When the diagnostics tests in ROM have been carried out, the BIOS looks for a disk that contains the very first part of the operating system, the boot sector. Using the very basic information in the boot sector, the computer can locate the rest of the operating system instructions on the disk, then read them into RAM. The operating system thus takes over control of the system. In a basic PC configuration, this leads you to the DOS prompt, an example of which is shown in Figure 2.7.

Most PC operating systems use a simple alphabetical convention to identify disk drives, starting with A for the first floppy disk drive, B for the second floppy if one is installed, then C for the first hard disk, D for the second, and so on. The prompt in Figure 2.7 suggests that the BIOS found the operating system on drive C, although this is not necessarily the first place that it looked.

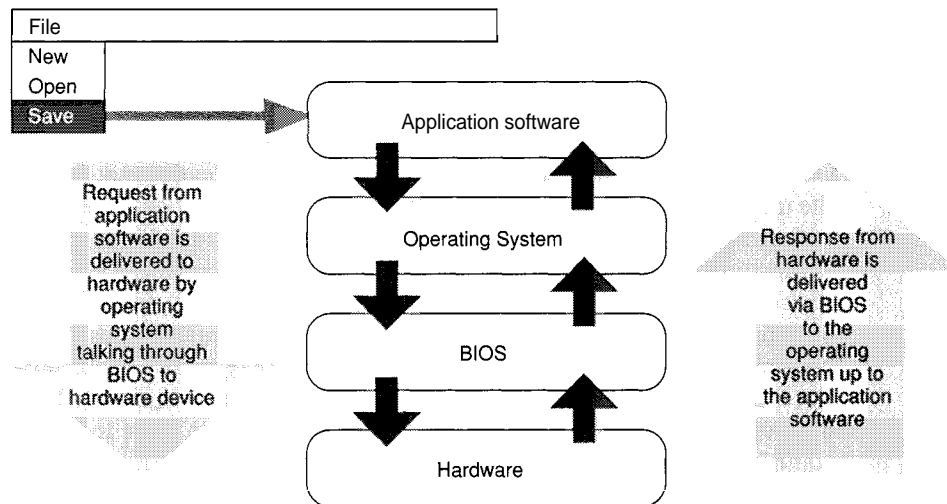


Figure 2.6 The relationship between BIOS, OS, data, and applications.

Traditionally, the BIOS looks to drive A first when it is attempting to locate the operating system, only checking drive C if there is no floppy disk in drive A or the disk in drive A lacks the operating system files. This process is called *seeking*. More recent versions of BIOS allow you to change the seek order. This change can be made through a program called SETUP, which is part of the BIOS program stored in ROM.

Typically you access SETUP by turning on the computer, then pressing the DEL key, or a combination of keys, when prompted to do so. The SETUP program allows you to adjust many different aspects of hardware operation, storing the results in

```
Microsoft (R) Mouse Driver Version 8.20
Copyright (C) Microsoft Corp. 1983-1992. All rights reserved.
Mouse driver installed
Stacker 4.00 for Windows & DOS (c) 1990-94 Stac Electronics, Carlsbad, CA
Registered to:
    Stephen Cobb
    Cobb Associates

C:\>
```

Figure 2.7 The DOS prompt.

CMOS. As you can see from the "Extended setup" screen in Figure 2.8, you can even disable "floppy drive seek at boot." This means that the system would only look for the operating system on the hard disk. I will discuss the value of this feature later on.

Prompt service

Returning to the DOS prompt, what are the four characters in Figure 2.7 telling you? The C indicates that the currently active drive is drive C. For example, if you type DIR and press the Enter at this point the directory command will list the files on drive C. In other words, drive C is currently the *default* drive. ("Default" is one of those words that computer people have taken over. In this context, it means "what the computer assumes unless you tell it otherwise.") Drive letters are customarily followed by a colon to distinguish them from programs (thus you can enter A: to select drive A, whereas entering A on its own tells DOS to execute a program called A.COM, A.EXE, or A.BAT).

While C: indicates the current drive, the backward slash in the DOS prompt indicates the current directory: the root directory. Directories form part of a hierarchical filing system, branching out like the roots of a tree, as can be seen in Figure 2.9. The directories that are created below another directory are referred to as *sub-directories*. The location of a file within the tree is referred to as its *path*, as in C:\123\WORK, which is the path to the file BUDGET.WK3.

The flashing thing at the end of the line is the cursor. This is where the next thing that you type will appear. In fact, until you type the name of a command that DOS recognizes, then press the Enter key, the PC does nothing.

While many of us are familiar with the DOS prompt as it appears in Figure 2.7, this is not the default DOS prompt. The prompt that you get when you first install DOS

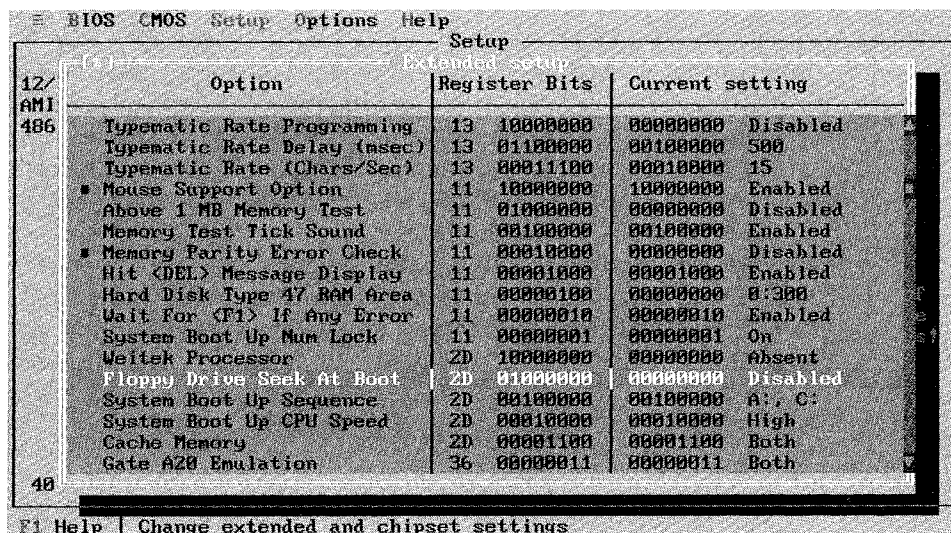


Figure 2.8 Viewing extended BIOS options in SETUP

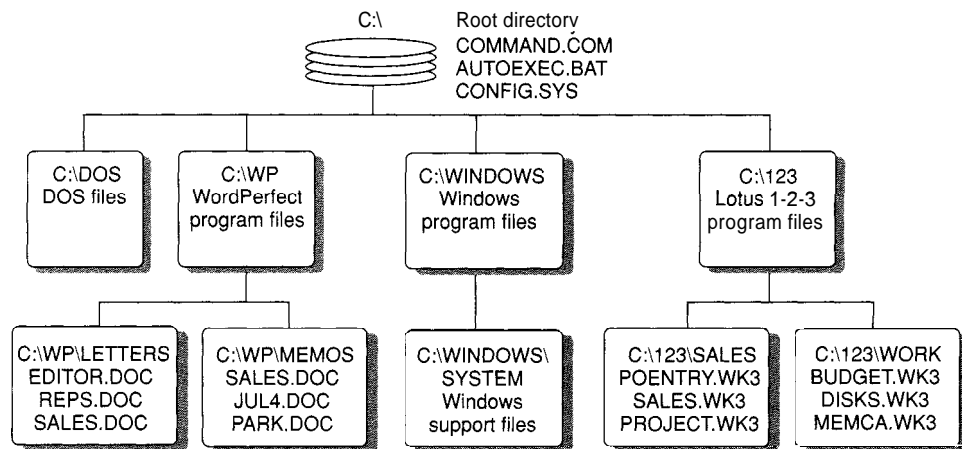


Figure 2.9 A typical DOS directory tree

is less informative. It only tells you the drive letter (appearing as simply `C>`). It is customary to use the `PROMPT` command to improve on the basic DOS prompt. The `C:\>` prompt is produced by entering the following command:

```
PROMPT $P$G
```

This tells DOS to display the current path (`$P`) followed by the greater than sign (`$G`), which results in `C:\>` when the current path is the root directory (`\`). If you change directory (for example, to the 123 directory on drive C), then the prompt would reflect that, as in `C:\123>`. The following command tells DOS to display the date as well as the path:

```
PROMPT $D $P$G
```

After you have issued the `PROMPT` command, the setting stays in effect until you turn off the machine or reboot it. If you enter `PROMPT` on its own without any arguments you will get the plain default prompt. So how do you make sure you get the more informative version of the prompt every time you boot up the system? This is handled by a special file `AUTOEXEC.BAT`, which I will discuss in detail in a moment.

In the case of the Macintosh, the loading of the operating system results in the drawing of the desktop, presenting your files and folders ready for you to begin work. The Macintosh has a filing system equivalent to that on a PC, using folders instead of directories, and folders within folders for subdirectories. The root directory on a Macintosh is the desktop itself. In Figure 2.10, you can see a Macintosh booted from the hard disk. You also can see a series of open file folders forming subdirectories.

System configuration

When the operating system is loaded into a personal computer's RAM, certain assumptions are made about how the operating system is to carry out its work. The assumptions are known as the *default settings*. You might want to tell the operating

system to use settings other than the defaults. A typical example is the FILES setting on a DOS system. The FILES setting determines the total number of files that the system can open at once. The default setting is 2. Many application programs require a much larger number for this setting, such as 40. Consequently, you need to tell DOS to deviate from the default.

This is accomplished by the entries in a special file called CONFIG.SYS, which resides on the boot disk. This file is simply a list of statements, each on a separate line, that tell DOS how you want your system to be configured. For example, this line tells DOS to change the FILES setting to 40:

```
FILES=40
```

A typical CONFIG.SYS file will contain several more statements to set up the system in a specific configuration. For example, DOS might need to be told how to handle memory as well as disk drives and other devices that attach to your PC. This is done through files called *device drivers*. A device driver contains a set of information used by DOS to relate to the device. The following CONFIG.SYS entry is used to load a driver for a special adapter (known as SCSI for Small Computer System Interface) to which several storage devices might be attached:

```
DEVICE=C:\CORELDRV\FDSCSI.SYS
```

Operating systems other than DOS use roughly comparable methods to record user preferences and system configuration. For example, on the Macintosh, the user's choice of desktop pattern is recorded in the System Folder. This preference actually is stored on disk when the Mac is shut down. Unlike most DOS systems, the Macintosh has a formal procedure that needs to be followed when the machine is




Figure 2.10 A typical Macintosh desktop.

turned off. Part of this procedure is the updating of hidden System files to record some of the preferences current at the end of the session. Many Macintosh applications also record user preferences in files, some of which also are placed in the System folder.

Not all Macintosh and PC system preferences are stored in configuration files. The PC uses CMOS, as described earlier, and the Mac stores such parameters as RAM cache size in something called PRAM, not a perambulator, but a piece of RAM known as Parameter RAM. Unlike regular RAM on a Mac, PRAM is maintained by the battery, like CMOS in a PC, allowing it to retain information from one session to the next, regardless of which disk you use to boot the system.

The System folder on a Macintosh also contains driver files for special hardware, such as the LaserWriter. These driver files usually are listed as Chooser documents when you view the System folder by Name. By including these drivers in the System folder, you ensure that they are available from the Chooser during operation. Mac System folders often contain two other types of files that are important to system startup: CDEVs and INITs (although INITs now are referred to as Extensions). Both of these are programs, which are described in the next section, that make useful features available automatically.

From a security perspective, the value of configuration files is that they typically are read into memory right after the operating system and before any programs can be run. By referencing some sort of security device in the configuration file, you can exercise some control over system access. You also can control access by keeping the required configuration file separate from the rest of the system. Suggestions on how to do this will be made later.

First programs

When the operating system and configuration information have been loaded, your personal computer is apt to just sit there, waiting for you to tell it what to do. The exception to this is the use of autoexecuting features. On the PC, autoexecuting is achieved through a specially named file, AUTOEXEC.BAT, that the operating system always looks for when booting up. Within this file is a series of instructions, telling the operating system to run a program or execute a command. A simple AUTOEXEC.BAT might look like this:

```
WECHO OFF -- Stops screen repetition of commands.
PROMPT $D $P$G -- Customizes the normal DOS prompt.
PATH C:\DOS;C:\WINDOWS -- Creates the PATH setting.
KEYB UK -- Runs the KEYB program for a UK keyboard.
```

For the most part, the instructions in AUTOEXEC.BAT are ones that you would otherwise have to type from the keyboard at the DOS prompt. They are placed one per line with the end of the line acting like the Enter or Return key used to enter the instructions. You can use AUTOEXEC.BAT to set a special prompt, set a path, or run a program. For example, on a PC dedicated to run payroll worksheets, the AUTOEXEC.BAT can load 1-2-3 automatically as soon as the system is turned on. Because there are autoexecute features in 1-2-3 and many other programs, including Windows, you thus

have the ability to load program and data files automatically, permitting a completely automated procedure from power on to the point of data entry.



More Bat Tails

As you might know, AUTOEXEC.BAT is a batch file, a special type of file used to execute a series of commands. You will find an introduction to writing batch files in appendix B.

To launch an application at system startup on a Macintosh, you used to employ the special option on the Macintosh menu (select the application, choose Set startup from the Special menu selection and check the Application box). This now has been replaced by the Startup folder, which is placed within the System folder. Applications to be launched after the system has been loaded are placed in the Startup folder, much like the Startup program group in Microsoft Windows, which launches applications after Windows has been loaded.

More extensive automation features are available on the Macintosh if you use Apple Script. While the Mac operating system does not have a direct equivalent of an AUTOEXEC.BAT, or of batch files in general, you can use Apple Script to execute any sequence of commands. This is similar to a macro system. *Macros* are collections of commands and operations that typically are carried out with a single keystroke (macro implying many keys under one). For example, the commercial applications AutoMac and QuickKeys allow you to create macros that executes when your Mac starts.

Another aspect of autoloading on the Mac is the CDEV and INIT/Extension files, which, when correctly installed in the System folder, cause software to be loaded automatically when the system starts up.

A typical CDEV is Monitor, which is used to allow switching between display modes. Placing the Monitor file in the System folder allows the Monitor icon to appear in the Control Panel. Such files are referred to as *Control panel documents*.

A typical INIT/Extension is QuickTime, which allows you to place QuickTime video clips. Another is the After Dark screen saver. Placing an Extension file in the System folder means that the program automatically is loaded into memory as the system starts up. This is equivalent to a RAM-resident utility on a DOS machine, otherwise known as a TSR program (for *terminate-and-stay-resident*).

Autoexecute Override

The ability to load programs automatically upon booting is a valuable one for programmers seeking to provide security features missing from the popular personal computer operating systems. However, it is important to bear in mind that autoexecuting features usually can be overridden. If you turn on a Macintosh while holding down the Option key, the Shift key, or both, some Extensions will be deactivated, or you will have the option to continue startup without them. If you press Ctrl-Break while starting up a PC that has an AUTOEXEC.BAT file, you will get this message:

```
Terminate batch job (Y/N)?
```

You can type Y and press Enter to circumvent the instructions in AUTOEXEC.BAT. Newer versions of DOS allow you to completely or selectively override both CONFIG.SYS and AUTOEXEC.BAT. When you boot up a PC that is running a recent version of DOS, you typically see a message *StartingMS-DOS* appear on the screen for about two seconds. This allows you time to press the F5 or F8 keys to interrupt the loading of the system.

To completely skip the CONFIG.SYS and AUTOEXEC.BAT instructions, you press F5. Alternatively, you can press F8 to selectively process CONFIG.SYS and AUTOEXEC.BAT instructions, one line at a time. This startup-interrupt feature was introduced to help users diagnose configuration problems. Different device drivers that might be in conflict can be selectively disabled. However, this also means that items in CONFIG.SYS and AUTOEXEC.BAT that provide security features can be disabled. For this reason, there is a command, called SWITCHES, that can be installed in CONFIG.SYS to block the F5 and F8 keys. The following line will disable the F5 and F8 and skip the two-second *StartingMS-DOS* delay:

```
SWITCHES= /N /F
```

Secure Boots

Having seen what happens when a personal computer is turned on, you can begin to see some of the security aspects of the process and the possibilities for preventing unauthorized access.

The floppy boot

Although they are slower and have less storage capacity than machines with hard disks, personal computers that rely on floppy disks are inherently more secure. This is because any unauthorized use requires a disk, and disks are easy to lock up. If you are running a personal computer that has no hard disk, someone still can boot up the system using their own copy of the system files; however, if the floppies holding your application and data files are locked away, there is not much point in an unauthorized user starting up your system, except to put in time on it.

The unbootable hard disk

Today's operating systems and applications are so large that a hard disk is almost a necessity. One of the first tasks when setting up a hard disk computer is placing the system files on the hard disk. When you are using PC-DOS, these files are called IBM.BIO.COM, IBMDOS.COM, and COMMAND.COM. If you are using MS-DOS, the first two files are called IO.SYS and MSDOS.SYS, respectively. While you probably are familiar with seeing COMMAND.COM in lists of files, you might not have seen the first two files listed. That is because DOS hides them. (There will be more later about hidden files.)

Working together, these three files enable the computer to be booted by the hard disk, which is the fastest and easiest way to start the system. However, you do not have to place the system files on the hard disk. If you remove or disable the system

files, you still can boot the computer from a floppy disk, then use the hard disk. This means that a person who does not have a system boot disk will have a hard time getting to the hard disk on the computer. This is one way to protect the information on the hard disk from casual interlopers.

Disk debooting

The procedure for removing the system files from a hard disk should only be followed after creating a floppy boot disk, then only by those experienced with DOS. There are two levels of boot-proofing. To get to the first level, you delete or rename `COMMAND.COM`. If you delete `COMMAND.COM` from the boot disk, then DOS cannot find it and cannot complete the boot process. Renaming is as effective and can be used in batch files that temporarily deboot a disk.

The second level of debooting involves deleting the system files. These are called `IBMBIO.COM` and `IBMDOS.COM` if you are using PC-DOS and `IO.SYS` and `MSDOS.SYS` if you are using MS-DOS. These files normally are hidden so you will get the response `File not found` if you simply enter:

```
DEL IO.SYS
```

You need to unhide the files before you can delete them. To do this, you can use the `ATTRIB` program that comes with DOS or use a utility such as `FA`, which comes with Norton Utilities. These programs allow you to alter what are called *file attributes*. Each DOS file has four attributes: read-only, archive, system, and hidden. The read-only attribute can be used to prevent important files from being erased or altered. The archive attribute shows whether a file has been backed up or not and is used by programs that create incremental backups—that is, backing up only those files that have been changed since the last backup.

The system attribute is assigned to files used by the operating system, and the hidden attribute is used to make files invisible to the `DIR` command. Using `ATTRIB`, you can display file attributes and change them, as shown in Figure 2.11. Most of the files on the disk have the Archive status, meaning that they have not yet been backed up and need to be archived. The first two files on the list carry hidden, system, and read-only status. These two files are the part of DOS that you do not see during normal operations.

As well as showing you the attributes of your files, `ATTRIB` allows you to alter them. For example, you can take the hidden, system, and read-only attributes off the system files, as shown in Figure 2.11. This allows you to see the files as part of a regular directory and also allows you to delete them.

Once the system files have been removed from a disk, that disk can no longer be used for booting the system, even if `COMMAND.COM` is on the disk. Thus, it is not necessary to remove `COMMAND.COM` to make a disk unbootable, which is an important point that will be addressed later.

A floppy boot disk

To create a floppy disk that will properly boot a hard disk system, you will have to consider what files the system needs in order to operate. In the case of DOS-based ma-

```

C:\>ATTRIB A:
      SHR      A:\IO.SYS
      SHR      A:\MSDOS.SYS
      SHR      A:\DBLSPACE.BIN
      R        A:\COMMAND.COM
A      A:\AUTOEXEC.BAT
A      A:\CONFIG.SYS
A      A:\FILES.DAT
A      A:\IDE.DSK
A      A:\INSTALL.NLU
A      A:\ISADISK.DSK
A      A:\NE2000.LAN
A      A:\SERVER.EXE

C:\>ATTRIB -H -S A:\*.SYS

C:\>ATTRIB A:\*.SYS
      R      A:\IO.SYS
      R      A:\MSDOS.SYS
A      A:\CONFIG.SYS

C:\>_

```

Figure 2.11 An example of using the ATTRIB command

chines, in addition to the system files (IO.SYS, MSDOS.SYS, and COMMAND.COM), you probably will want CONFIGSYS and AUTOEXEC.BAT in order to give preliminary instructions to the system.

Without a correct CONFIG.SYS file on the floppy boot disk, problems in running the system can arise, preventing the interloper with a copy of the DOS disk from getting very far. For example, the correct CONFIGSYS might contain a DEVICE= statement that refers to a device driver, without which the system will not operate properly. The following line tells DOS how to manage memory, and without a CONFIGSYS file on the boot disk that contains this line, you will have a hard time running Windows on that machine:

```
DEVICE=HIMEM.SYS
```

In order for your floppy boot disk to work effectively, you need a copy of the CONFIG.SYS file that normally would be on the hard drive to the floppy boot disk. You should copy to the floppy the necessary SYS files, such as HIMEM.SYS, to which CONFIGSYS refers. Technically, you can refer to drive C for these SYS files, but it is safer to place them on the floppy disk, with local references (such as DEVICE=FDSCSI.SYS rather than DEVICE=C:\CORELDRV\FDSCSI.SYS). To make the system more secure, you then can remove the SYS files from the hard disk (after first making sure that you have several backups stored away, just in case). For example, if your system uses a SCSI hard disk that depends on the line DEVICE=FDSCSI.SYS and you then move FDSCSI.SYS to the floppy boot disk, you make it very hard for anyone to access that SCSI drive without the floppy boot disk.

Paths and programs

The role of AUTOEXEC.BAT also is important when booting a DOS system. Many AUTOEXEC.BAT files set up the appropriate environment for the applications on the hard disk, including SET and PATH statements. If you boot from a floppy without the correct AUTOEXEC.BAT, then running some applications can be difficult. For example, the PATH statement tells DOS which directories to look in when you ask it to run a program.

Take the case of a DOS utility program like CHKDSK, which reviews the use of space on the disk. This actually is a program file called CHKDSK.COM. To use the program, you enter CHKDSK at the DOS prompt. Suppose that you do this while the root directory is the current directory. The first thing DOS does is look in the root directory for the file CHKDSK.COM. Whenever you ask DOS to run a program from disk, it always looks in the current directory first. If the program is not in the current directory, then DOS checks to see if there is a PATH setting. If there is not, you get a message like this:

```
Bad program or file name
```

Suppose that there is a PATH setting, such as PATH=C:\DOS. This tells DOS that, if the program you request is not in the current directory, it should look in the DOS subdirectory of drive C. If the file is in C:\DOS, then the program is executed. Note that you could run CHKDSK from the root directory without a PATH statement by entering \DOS\CHKDSK. Alternatively, you first could make DOS the current directory by entering CD\DOS, then run the program by entering CHKDSK.

To determine the current PATH setting on a PC, you simply enter PATH at the DOS prompt. To set a PATH, enter PATH= followed by a list of directories, separated by semicolons, as in:

```
PATH=C:\;C:\DOS;C:\123
```

This tells DOS to look in the root directory of C as well as the DOS and 123 directories. Entering a PATH= statement overrides previous PATH settings. The AUTOEXEC.BAT on a boot floppy can be the same one used when booting from the hard disk; however, you might want to customize it in order to make sure that drive C is the active drive by the time the instructions in AUTOEXEC.BAT are completed. Adding a line that simply contains C: will accomplish this.

Bootless problems

One drawback to a floppy boot system on DOS systems is that DOS will continue to look for COMMAND.COM on the drive from which you booted. You might wonder why DOS would need to look for COMMAND.COM after reading it from disk in the initial booting procedure. After all, COMMAND.COM is the operating system's command interpreter, loaded into RAM for the duration of the session. Unfortunately, it is not quite as simple as that.

Part of COMMAND.COM is always resident in memory, and part is transient. The transient part handles your command requests from the prompt line and is not needed when you are running an application. To give you more memory when running applications, the transient part is unloaded, then reloaded from the COMMAND.COM file when you quit the application. If you boot from a floppy in drive A, load WordPerfect, take out the boot disk (perhaps to insert a data disk), then exit back to DOS, you might get a message like this:

```
Insert disk with COMMAND.COM in drive A, and press any key to continue
```

DOS is simply looking for COMMAND.COM to reread the transient portion into memory, and DOS assumes that COMMAND.COM is on your boot disk. You can get around this problem in several ways. One technique is to copy COMMAND.COM onto a RAM disk, then use a SET statement in the AUTOEXEC.BAT file to tell DOS that the RAM drive is the place to look for COMMAND.COM. For example, the following lines in an AUTOEXEC.BAT will do the trick:

```
COPY A:\COMMAND.COM D:\
SET COMSPEC = D:\COMMAND.COM
```

In this case, D is the RAM drive. (The letter assigned to the RAM drive is the next one after the last letter used by the other drives in your system. If you have two hard disks, C and D, then the RAM disk will be E.) The COMSPEC setting determines where DOS looks for COMMAND.COM. You can set up a 256K RAM drive in extended memory with the following line in your CONFIG.SYS:

```
DEVICE = C:\DOS\RAMDRIVE.SYS 256 /e
```

Another technique to get around the floppy boot "missing" COMMAND.COM problem is to leave COMMAND.COM on the hard disk, then use a SET statement to tell DOS that the hard drive is the place to look for COMMAND.COM, as in:

```
SET COMSPEC = C:\COMMAND.COM
```

You might wonder how you can leave COMMAND.COM on a disk when the disk is not to be used for booting. Remember that DOS requires two other files (IO.SYS and MSDOS.SYS) besides COMMAND.COM in order to boot, as described earlier.

Special cases

A couple of situations deserve special mention when it comes to keeping people out of high-capacity storage devices. Some hard disks are removable, working much like big floppy disks. Prime examples are IOMEGA Bernoullis, the Syquest cartridge systems, and optical disk systems. These devices offer a good compromise between the high speed and large capacity of a conventional hard drive and the removability of a floppy with all its inherent security advantages. Slip your cartridge out of the drive and into a locked drawer and your data is very safe. In Figure 2.12, you can see some of the cartridges used in these systems.

You can even remove a cartridge while your computer still is running, so secure lunch breaks are a possibility. Just save your work, eject the cartridge, and put it




Figure 2.12 Removable cartridge drive media.

away. Obviously, you should not take out the cartridge while an application actually is using it. However, as long as you are not in the middle of an operation that requires cartridge access, you can remove the cartridge without having to quit the program. If someone tries to use the program while you are away, they will get an "error reading drive" message.

You can install these systems on both Macs and PCs. The cartridge can be set up as the boot device, or you can boot from a different hard disk or a floppy, then access the cartridge. For the operating system to recognize the cartridge, a device driver is needed. This means an added measure of security is possible. If you configure the system so that the cartridge is not a boot device, you can boot from a floppy on which the device driver has been placed. If you then lock away the floppy, you can prevent people using the system, even if they find a cartridge lying around.

Other removable storage systems include complete hard disks in a special case, such as the Quantum Passport series. As hard disks continue to shrink in physical size and expand in storage capacity, removable units become increasingly feasible. Hard drives that attach to your PC via the parallel printer port offer another approach to "removable" storage.

Later chapters **will** return to the subject of removable storage media and system access controls. The purpose of the last few sections has been to clarify the elements of desktop systems and to show where some of the security issues arise. Knowledge of these elements will be required in several other areas, such as viruses, which often infect computers during startup, and password controls, which

sometimes are implemented through system and configuration files. For example, the Lock access control program shown in Figure 2.13 is installed as a device driver in CONFIG.SYS.

Under Lock and Key

So far in this chapter, you have seen that a knowledge of system components is important if you are to understand security threats and remedies. I also have stressed that users of personal computers need to be more diligent in backing up their data and more vigilant in general. The next item on the list of security basics is physical security, as in theft deterrence. In other words, making sure that your computer stays where you put it and is there when you need it. There are several levels at which a lock and key can work for you. In the following sections, some straightforward but often overlooked security measures are reviewed.



Hasty Gangs

Thirteen gang members allegedly on their way to steal computer parts in Sunnyvale were intercepted on the freeway by a phalanx of squad cars early today. Police said the "hasty gang"—a group organized for a single crime—had planned to make off with more than \$1 million worth of parts from a warehouse belonging to Amdahl Corporation.

San Jose Mercury News, January, 1993

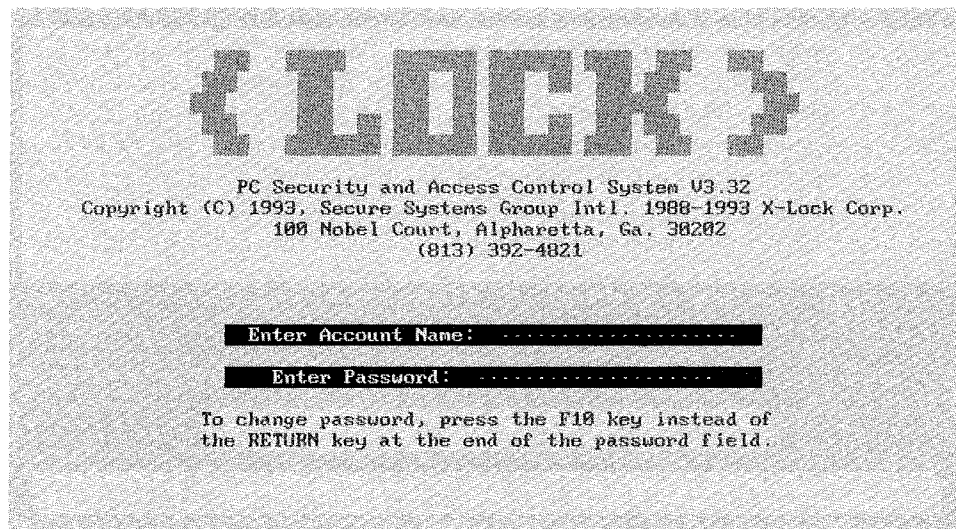


Figure 2.13 The Lock access control program

Secure premises

Access to rooms containing valuable computer equipment and information should be controlled, even during business hours. Locking the door might seem like a pretty obvious precaution, but the obvious is worth mentioning for several reasons. For a start, the obvious is all too easily overlooked. For example, many companies have a personal computer in their reception area so that, between answering phones and greeting visitors, the receptionist can do word processing and data entry. Several risks are involved in this arrangement.


The presence of the personal computer alerts all comers to the fact that the office is computerized. There might only be a short distance from the reception area to the parking area, increasing the temptation for a petty thief. In other words, try to keep general security principles in mind when locating personal computer equipment, as illustrated in Figure 2.14. This includes making sure that reception areas are always attended. Failure to greet visitors is not only rude, leaving unannounced visitors with an unguarded personal computer is asking for trouble. Several minutes is all it takes to load a personal computer into a van and drive off.

If you have to leave equipment in an area to which the public has unsupervised access, make sure it is bolted down (systems for bolting down can be found in chapter 4). Another reason for stressing obvious precautions, such as not placing personal computers in front of unguarded street level windows, is that the personal computer has become both a commodity item with good street value and an almost invisible part of the office furniture. Many of today's office workers are so used to seeing personal computers around that they tend to overlook their value. (I have seen office staff carefully lock up transistor radios at the end of the day and yet leave floppy disks on the desk and give no thought to locking up the computers.)

Occasional reminders from management about correct procedures and the need to protect the organization's personal computer investment are a valuable antidote to complacency and the tendency to take personal computers for granted. Note the word "occasional." Bear in mind that reminders made too frequently are themselves likely to be taken for granted. If you are charged with issuing security proclamations, try to make them interesting, and thus more likely to be noticed. See the example in Figure 2.15 for some ideas.

One factor that makes the more obvious security practices more important these days is the rapidly decreasing size of personal computers. Of course, as some notebook manufacturers have been quick to point out, the ability to slip an entire PC into a lockable desk drawer is a security advantage. However, the shrinking size of technology can make matters easier for the would-be thief, providing more value per easily lifted pound.

In the days when computers could be measured in cubic yards rather than cubic inches, they were housed in purpose-built rooms that had climate as well as access controls. It was natural and fairly straightforward to secure these computer rooms. Today's personal computers can be as powerful as machines many times their size from just a few years ago, but they do not require special rooms. This encourages our tendency to think of personal computers as being akin to telephones or photocopiers. However, the consequences of losing a personal computer are potentially far more damaging.



- A. Reception monitored and staffed at all times, controls access to rest of site, card id on inner door.
- B. Note poor location of computer in front of window.
- C. Better location of computer system away from window. Note closed door.
- D. Computer location away from window good but watch for open doors.
- E. Poor location next to window obscured from view by bushes, a prime target for burglary.
- F. Warehouse area with loading door needs vigilance. Do not leave unstaffed when doors are open.
- G. Open plan work area—hardware should be secured to furniture and system access controls used.
- H. Equipment is kept well clear of window for better security.

Figure 2.14 Tips for locating and securing computer equipment.



Figure 2.15 Security warnings created with a desktop publishing program



RAM Raiders

The term Ram Raid originally was coined by the press to describe a new tactic among thieves—that is, breaking into a bank or shop by literally ramming a vehicle into it. However, the term has taken on another meaning as crooks have discovered the tremendous value-to-weight ratio of memory chips and other computer components, which can exceed that of gold and platinum.

In Silicon Valley in the early 1990s, several high-profile attacks were made on chip fabrication plants and trucks used for chip deliveries. Two of the big attractions of memory chips are the ability to sell them just about anywhere in the world and the lack of serial numbers, which made them virtually untraceable. When chip makers responded by adding serial numbers, there was an upsurge in chip thefts from existing computers. For example, according to Secure Computing, in January of 1995, “thieves stole £25,000 worth of chips from British Telecom’s offices in High Holborn.” One month earlier, the Automobile Association “suffered an even worse disaster, losing £125,000 worth in just one night. As well as causing considerable disruption through the loss of the use of the hardware, the physical disruption of disemboweled computers was devastating.”

Locked cubicles

Given the open plan structure of many of today's offices you might be saying, “It's all very well to talk of locking the office door, but I don't have one.” In such situations, you might need to resort to different tactics. You can buy all manner of devices for fixing personal computers to desks and otherwise locking them away, and a representative sample is given in chapter 4.

Typically, the responsibility for general security in an open-plan office rests with management. Various strategies are used, the most common of which is restricting access to the office. This can be done through a front desk/back office layout, a card or I.D. badge reader, or a combination lock, to name just a few of the approaches. It is worth noting the generally accepted principle that organizations have a responsibility to provide a secure work environment for employees.

Locked computers

When IBM introduced the PC AT in 1984, much ado was made about the fact that it came with a lock on the front that disabled the keyboard. Some members of the computer press hailed the “key to the AT” as a new workplace status symbol, right up there with the key to the executive washroom. Those who were lucky enough to get an AT on their desk could lock out less fortunate users. These days most PCs have keyboard locks on the front, similar to the one shown in Figure 2.16.

After the initial hype, the reality of day-to-day operations sets in, and the keys were forgotten by many users. In fact, these days many people would have difficulty finding the keys. Other people can show you the keys still sitting in the computer. Both situations pose a security threat. Yet proper use of these keys provides one of the most effective low-level deterrents to unauthorized access. If you turn the key to




Figure 2.16 The keyboard lock on Digital's Venturis 590 PC

the locked position and remove it, only a determined interloper can use your computer. Yes, you can bypass the lock by messing with the system wiring, but this usually involves opening the case of the system unit.

In the well-organized office, a key distribution and control system should be implemented to manage the issue, use, and return of keys. Each user might be issued a key while the supervisor retains the copy. When an employee leaves, the key should be surrendered. For more on physical key management, see chapter 6.

Locked up backup

By now, the security role of backup copies of important data and programs should be clear. Even if your computer system is stolen, you can get back to work fairly quickly if you still have backups of your applications and data. However, you cannot rely on backups if they are treated carelessly. The backup media, be it disks or tapes, should at least be locked away at the end of the day. Most offices have drawers or cabinets that can be locked. The casual burglar is more interested in hardware than data and is not likely to break into cabinets just for some extra tapes or disks.

One step up in backup security is a safe, preferably a fireproof model. This will give your data added protection against one of the worst natural disasters an office can experience, but even more secure is offsite storage of backup. Consider the construction company that does important accounting on a personal computer located in a portable office at a job site. Fearing vandalism, the site manager takes the backup tape home every night. This protects the backup from the perils that might befall the site office. Of course, it is possible that the tape might be damaged or lost

while in the care of the site manager, but this is the backup. For the company to experience data loss, both the original data at the site office and the copy carried by the site manager would need to be damaged. If the data you are backing up is mission critical, then offsite storage is highly recommended. (See chapter 8 for more about offsite storage options.)

Removable disks and computers

There are definitely tradeoffs between component size and security. A very small computer with large storage capacity offers the chance to lock the whole system away in a desk drawer but also makes it easier for the whole system to be carried off under someone's arm. Storing a lot of valuable data on small, easily pocketed media presents the same dilemma. As storage capacities per cubic inch increase, there is less need to split information into small sections.

When the IBM PC first came out, a backup copy of a document like the manuscript for this book would have required at least half a dozen floppy disks. Now I can backup the text of dozens of books on a cartridge smaller than a book of matches. Generally this is very convenient, but consider the implication for company accounts. A high-capacity disk cartridge, like the ones shown earlier in Figure 2.12, can easily contain the entire accounting system for a medium-sized company.

On balance, the security benefits of easily removable high-capacity disks outweigh the downside. At least, this is what the FBI has concluded because it specifies removable hard disks for critical personal computer applications. I will return later to the advantages that such devices offer in terms both of backing up and of controlling access.

Basic File Protection

The terms *read* and *write* are used to signify the complimentary actions of getting information from a disk and placing it onto a disk. Writing to a disk applies to any kind of change to the disk, including erasing files from it. A simple way to prevent data on a disk from being lost or damaged is to use write-protection. You can write-protect your files either physically or through the operating system.

Physical write-protection

A file that is write-protected is referred to as a read-only file; it still can be read, but it cannot be deleted or even replaced by an updated version of the same file. The prime security benefit of write-protection is the prevention of accidental erasure of files. A reasonably competent user can defeat write-protection, but the innocent user is well-served by it. Physical write-protection is a standard feature on floppy disks, made possible by a small lever in the drive itself. On traditional 5.25" disks, the notch on the right of the disk jacket (shown in Figure 2.17) allows a drive to write to the disk.

A lever in the drive will move into the notch when an unprotected disk is inserted into the drive, allowing the disk to write to the disk. If you close off the notch with a piece of tape or one of the sticky tabs provided in packages of blank disks, then the disk is write-protected. This should be a normal precaution with archive copies.




Figure 2.17 Write protection on floppy disks

The newer 3.5" disks also feature write-protection, but in a slightly different manner. In the top right of the disk is a square hole, as shown in Figure 2.17. When you want to allow writing to the disk, this hole must be blocked, which normally is done by turning the disk over and sliding a rectangular piece of plastic across the opening. The drives that read these disks have a lever that allows writing to a disk only if it is kept out of the hole, somewhat the opposite of the system on larger drives. Some 3.5 disk drives use an optical system for write-protection. If the light passes through the hole, then writing is disabled.

Of course, write-protection only offers protection against accidental or casual interference, because the system is easily reversed. Furthermore, write-protecting does little help unless you also stick to the rules for disk handling. With careful handling, your disks will retain data for many years. See Figure 2.18 for the basic rules.

Not surprisingly, one of the most basic tips for older 5.25" drives is to place disks in sleeves when they are not in drives. This wards off fingerprints, coffee spills, and tobacco smoke, all of which pose a threat to the integrity of data on disks. The hard-jacketed 3.5" disks have built in shutters to cover the read/write access slot, but they should be shelved or otherwise removed from the work surface when not in drives (the disk jackets are by no means waterproof).

Software write-protection

When you store information in a file on disk, the operating system might store certain housekeeping information as well. This might include the date and time that the file was stored on disk or the name of the program that created the file. Some of this housekeeping information might determine whether or not the user can alter the file. For example, if you select a file on the Macintosh, then select **Get Info** from the File menu, you can check the Locked box, as seen in Figure 2.19.

This is the software equivalent of write-protecting the file, preventing files from being erased or altered. The file still can be used; however, if you read the file into memory and alter it, you cannot store the altered version over the locked original.

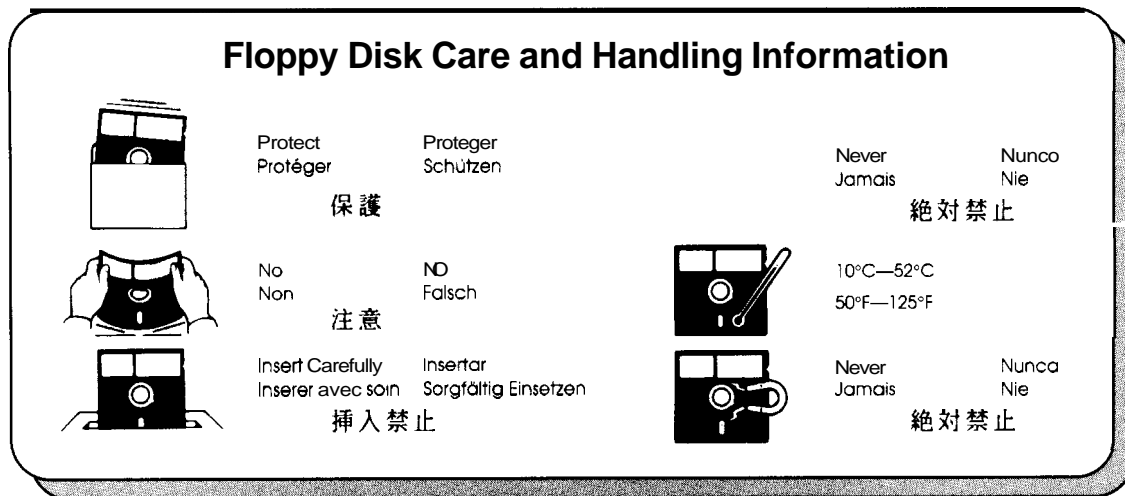


Figure 2.18 Rules for floppy disk handling

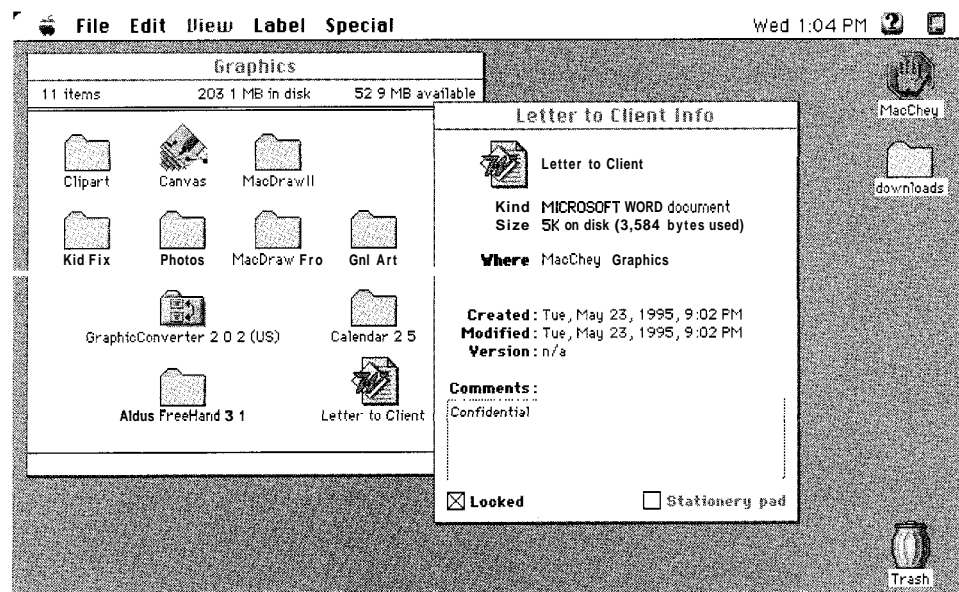


Figure 2.19 Locking or write protecting files on the Macintosh.

Locked Macintosh files are shown with the padlock icon in the By Names file listing. If you use the icon method of file listing on the Mac, you can detect protected files by clicking on the name. If the mouse pointer does not change to an edit cursor, then the file is locked. You can lock several files at once by selecting more than one file, with Shift-click, before selecting **Get Info**. The locked status of Mac files is preserved when they are copied from one disk to another.

On DOS systems, you can lock a file with the `ATTRIB +R` command, which activates the file's read-only attribute. In Figure 2.20, you can see how this works. You reverse the `ATTRIB` command by using a minus sign instead of a plus sign. Also note that you can alter the attributes of several files at once by using DOS wildcards. For example, to give read-only status to all of the program files in a directory that have the `.EXE` extension, you would enter:

```
ATTRIB +R *.EXE
```

This prevents accidental removal of these program files. The read-only status of DOS files is maintained when they are copied from one disk to another.

Note that locking files can create its own problems. Some programs write changes to the program file during use. If the file is locked, this is not possible, sometimes leading to unpredictable results. To make extensive changes to the DOS file attributes, you might want to employ a DOS utility program like QuickDOS, which presents a menu-driven file listing with a tag system in order to simplify selective attribute manipulation, as shown in Figure 2.21.

While You're Away from Me

So far you have seen some simple techniques for deterring people from stealing or starting up personal computers. However, this does not solve the problem of defending computers that are in the office, turned on, and unattended. This might be during bathroom or coffee breaks, over lunch, or during meetings.

```
C:\UTILS>ATTRIB +R *.EXE
C:\UTILS>ATTRIB *.EXE
R      C:\UTILS\ACUDUP.EXE
R      C:\UTILS\DISKFIX.EXE
R      C:\UTILS\DRUWTC.H.EXE
R      C:\UTILS\LE.EXE
R      C:\UTILS\LLPRO.EXE
R      C:\UTILS\MH-IDE.EXE
R      C:\UTILS\MONITOR.EXE
R      C:\UTILS\PSP.EXE
R      C:\UTILS\RAMTEST.EXE
R      C:\UTILS\SAVE.EXE
R      C:\UTILS\SETUP.EXE
R      C:\UTILS\SNOOPER.EXE
R      C:\UTILS\UNDEL.EXE
R      C:\UTILS\WINBENCH.EXE
R      C:\UTILS\WINP.EXE
R      C:\UTILS\WRLDTIME.EXE
R      C:\UTILS\ZIP2EXE.EXE
R      C:\UTILS\CARCH.EXE
R      C:\UTILS\XARC.EXE
C:\UTILS>
```

Figure 2.20 Using the `ATTRIB` command in DOS to protect files.

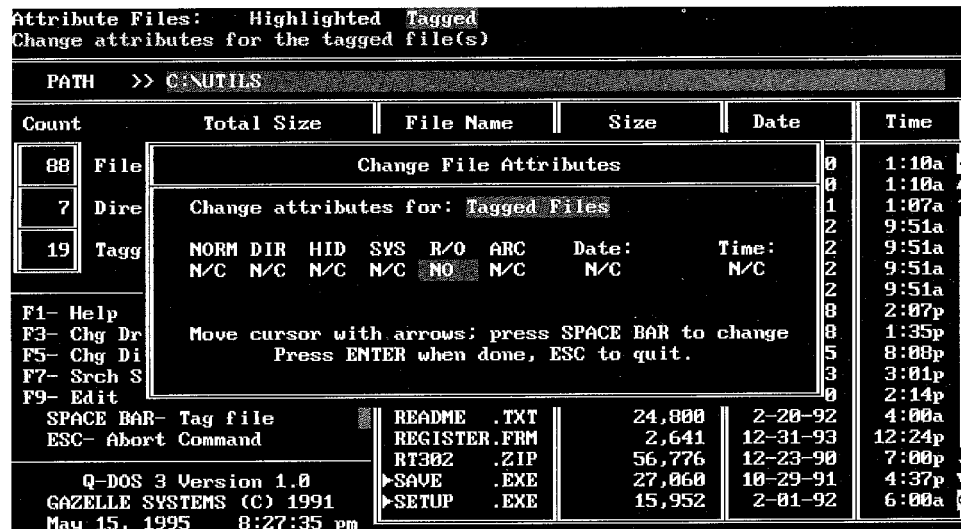


Figure 2.21 Changing attributes for a group of files with QuickDOS

See no evil

As a matter of policy, all users should always save their work before leaving their desks. This avoids all sorts of problems such as, "Who used my computer without saving what I was working on." However, turning off your computer every time you leave your desk is not only impractical, but it also can be bad for the computer. Apart from the fact that turning the system back on can take several annoying and unproductive minutes, it is possible that the hardware will last longer if you avoid, as much as possible, that initial rush of electricity through the circuits, the effort of bringing the hard disk up to speed, and the shock of power to the monitor.

On the other hand, leaving systems on for long periods involves few hazards. In most cases, so to speak, the only moving parts are the fan and the hard disk. While normal hard disks spin continuously, the delicate read/write mechanisms are at rest when the system is not receiving input. The one question concerning extended periods of personal computer use is heat damage. Any properly designed personal computer will maintain the temperature of components within working limits for extended periods. What can cause problems are variations in ambient temperature that exceed those limits. For example, a PC that works fine in the office during the day might have problems at night or on weekends when the normal air cooling systems are turned off to save energy.

A single PC can produce several hundred watts of heat, and even a small local area network can act like a two-kilowatt fan heater. Combine that with a hot day and a lack of air conditioning, and you could be looking at heat damage to sensitive circuit boards. Indeed, any time that the temperature in your office gets above 90°F (32°C), your personal computer equipment is liable to fail due to heat problems. Fortunately, recent advances in hardware design, such as those that comply with the Energy Star guidelines, have reduced electricity consumption and heat output. See chapter 5 for advice on preventative measures.

One source of avoidable wear and tear on unattended computers is the monitor. Some monitors are adversely affected by displaying the same image for prolonged periods of time. The image gets "burned in" to the phosphors painted on the inside of the screen. These phosphors normally glow when the electron beam hits them, producing the image that you see. In fact, they can glow very brightly, producing an image that is visible even in broad daylight. Try putting on a pair of dark glasses in the office. Your view of the rest of the office will be dimmed, but the screen will be perfectly legible. (For more about the physiological effects of using monitors, see chapter 5.)

Try this if you are in the office at night: Turn off the lights and turn on a personal computer that was used during the daylight, that is, one with the brightness turned up fairly high. Notice how painfully radiant the image is. If you leave it on for about five minutes, then turn it off, the image will appear to linger on the screen. On older monitors this actually caused physical damage (monitors that had been used to run Lotus 1-2-3 every day for years retained a legible image of the spreadsheet outline).

Locked out for lunch

The solution to screen burn-in is a screen blanker or screen saver, a program which, instead of sending an image to the screen sends either no image at all or one that changes to avoid "burn in." This also is the solution to the problem of preventing people from reading what you have on your screen while you are away from the computer. Microsoft Windows comes with a password-protected screen saver, the setup of which is shown in Figure 2.22. Password-protected screen savers for DOS (without Windows) and the Macintosh can be obtained at low or no charge. Examples are provided on the disk included with this book. The idea is to blank the screen on command or after a certain amount of time (set by the user) during which there is no mouse movement or keyboard input.

These days moving images are used more often than complete blanking because they let you know that the machine still is turned on (one problem with complete screen blanking is the risk that someone might assume a machine to be turned off just because there is no image on the monitor). The use of password protection on a screen saver means that, when the screen saver has been activated and someone presses a key or moves the mouse, a password is requested, as shown in Figure 2.23.

Failure to provide the correct password means that the screen stays locked. While the standard screen saver in Windows is not a particularly powerful access-control system (you usually can circumvent it by rebooting the computer), it is a useful defense against casual snoops.

Further Measures

To return to the theme of my opening remarks in this chapter, you actually can combine screen saving with security awareness. For example, the SAFEware security-awareness program comes with a screen saver that displays a series of security reminders, like the one shown in Figure 2.24. The slides change every few seconds, based on user settings.

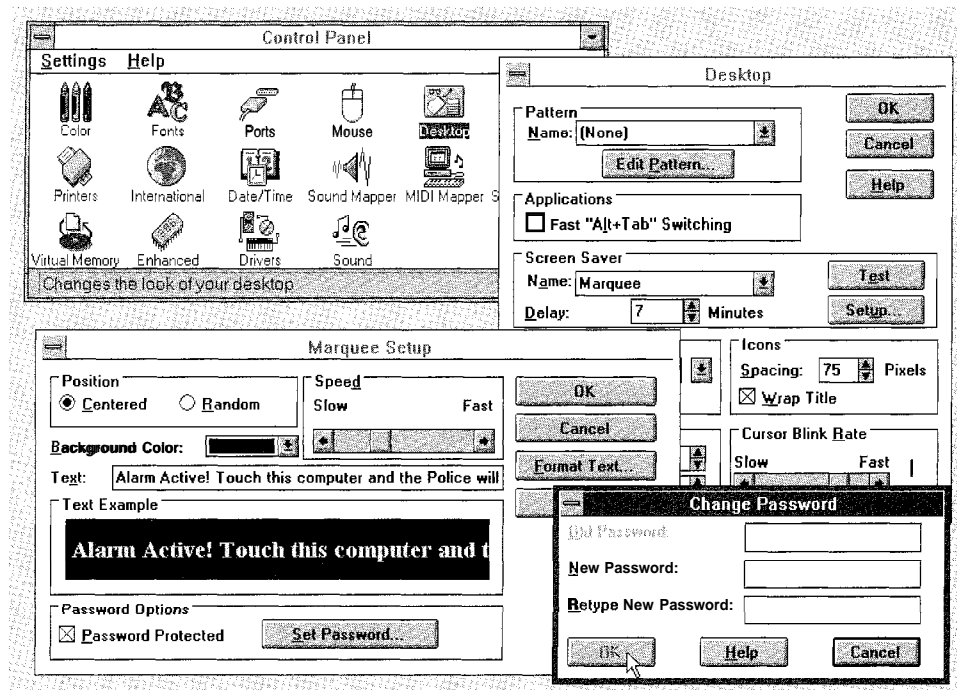


Figure 2.22 Setting up the Windows screen saver




Figure 2.23 Providing a password to unlock a screen saver

There are a number of other simple-but-effective security measures that can be instituted at low or no cost. These include password protection for files, which is provided by most major applications. Details on passwords for file access control can be found in chapter 7. The use of password schemes for system access control is covered in chapter 6. The use of maintenance tools, many of which are included with your system software, including disk defragmenters and optimizers, is covered in chapter 8. Advice on using anti-virus software is given in chapter 9.

The Network Connection

The basic elements of personal computer operation—such as the boot process and the interaction of BIOS, operating system, and application software—also apply to

network machines, whether they are workstations or file servers. Just as individual machines are controlled by an operating system, the operation of a network is controlled by a network operating system (NOS). The NOS might be added on top of a basic single-user operating system, such as DOS, creating an additional layer between the operating system and application software. This is the case with some peer-to-peer networks, such as Windows for Workgroups.

Alternatively, the NOS might be integrated into the operating system, giving it built-in networking, as in the case of Personal NetWare, which is part of Novell DOS 7, or AppleTalk, which is part of the Macintosh OS.

A third option is for the NOS to be an operating system in its own right, which is the case with some versions of Unix and NetWare. The concepts of file attributes and write protection also are used by network operating systems, which typically provide a greater range of attributes, sometimes referred to as *access* rights. The NOS normally will include software with which to manage access rights and attributes.

Security awareness among users of a network is perhaps even more important than among users of standalone machines, because networks are more likely to be relied upon for very important, if not mission critical, operations. As you will see in chapter 11, networking personal computers adds valuable control mechanisms to personal computers, but it also increases the number of fronts upon which an attack can be mounted. Thus, there is even more need for frequent, well-organized backup,

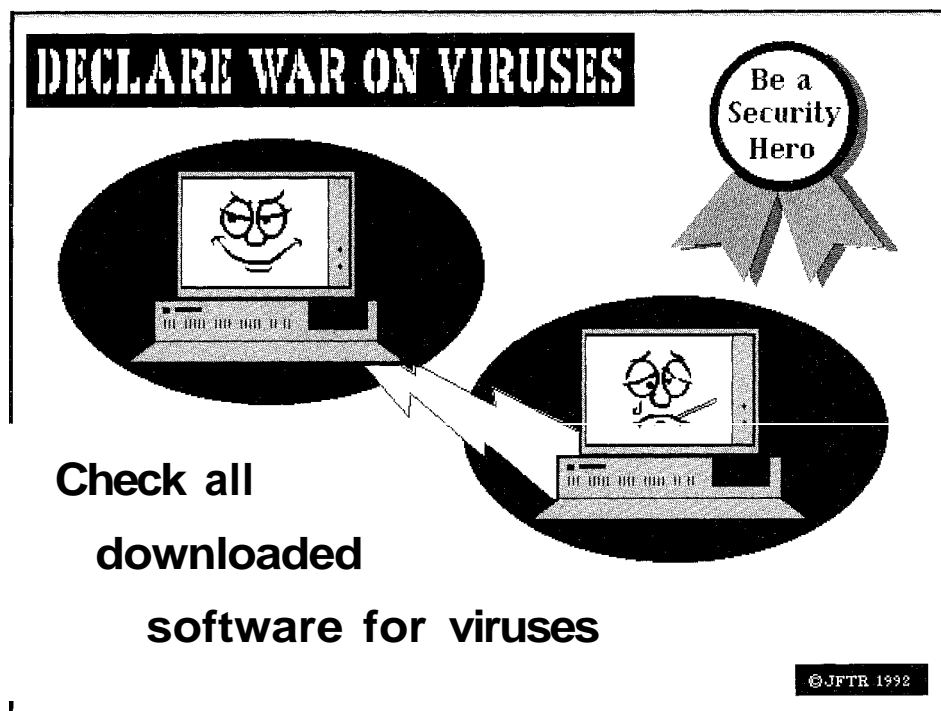


Figure 2.24 Security awareness can be combined with screen saving, such as this reminder from SAFEware.

which is made somewhat easier thanks to centralized storage of data on the file server. In turn, there is even greater need to physically protect the server, locating it in a separate room with strictly limited access if at all possible.

Summary


Raise awareness. Let everyone in your organization know that computer security is an important issue. Information is the lifeblood of the modern enterprise and job security depends upon defending it.

Backup is the first defense. Digital information is easy to copy, and backup copies are the first line of defense against theft, damage, errors, and omissions.

Understand your systems. You cannot defend what you don't understand. Learn how your systems work in order to anticipate threats, plug gaps, and implement defenses.

Don't overlook the obvious. Use the locks that are in place, manage the keys to all locks, and think like a thief to protect against theft.

Much can be done to achieve a basic level of security by using common sense, communication, and the system features already at your disposal. The next chapter looks at how you assess your risks, define security policies, and implement a security plan.



Security Planning Risk Analysis and Security Policy

W.E.B. connection: <http://www.nlsa.com/pclan/chap03.html>

*"There is always something to upset the most
careful of human calculations"*

SAIKAKU IHARA, 17TH CENTURY A.D.

*"The best laid schemes o' mice and men
Gang aft a-gley"*

ROBERT BURNS, 1785

This chapter helps you take some important first steps on the road to a secure computing environment. You will read how to assess the risks inherent in the use of personal computers and how to plan and implement appropriate defensive measures. The main focus of this chapter is planning within an organization; therefore, most remarks are addressed to support staff rather than SOHO (Small Office, Home Office) and group users. However, it would be helpful for all users to be familiar with the general principles of analysis, policy-making and implementation, that are presented here.

Questions, Concepts, and Cycles

Protecting personal computer-based information systems, from a multisite network to a handful of standalone PCs, requires answers to a number of questions. How do I put a value on information assets? How do I quantify the risks of attack? How do I weigh the potential losses from an attack against the cost of security measures? What kind of policies does my organization need to address computer security? What kind of advanced planning should we do to prepare for a major breach of security or a natural disaster such as earthquake, fire, or flood?

The analysis-audit feedback cycle

The previous questions are the main focus of this chapter. The answers will be embodied in a number of important documents. If it has not done so already, your organization should develop its own versions of these documents, which are listed in approximate chronological sequence:

Business Impact Analysis

Security Policy

Security Plan

Disaster Recovery Plan

Security Audits

The phrase "approximate chronological sequence" implies that the documents are developed one after another, each based on the preceding document. An exception is the Security Audits. These are comprehensive reviews of current levels of actual security. In other words, they describe what actually is happening as opposed to what should be happening, what users are doing rather than what they should be doing. Security audits can be performed at anytime. For example, they are used to assess the extent to which an organization has taken heed of the other documents, which is why they appear at the end of the previous list.

However, security audits also can be used to alert an organization to security problems, and thus begin the process of identifying risks and developing appropriate responses. However, there are many aspects of security that audits alone do not address, such as the likelihood of security weaknesses being exploited or the financial impact of such breaches. To answer these questions, you need a Business Impact Analysis. This is an assessment of potential threats, plus the consequences and likelihood of their materializing. The process of creating this document is known as risk analysis, which can be anything from a sole proprietor making notes about what could go wrong and what the cost would be to a major survey performed by a team of consultants. The end result in either case is a Business Impact Analysis that lays the groundwork for preparing sound security policies.

In the Security Policy, you state your organization's general commitment to security, then layout the rules and regulations on specific issues. These might include confidentiality, ownership of intellectual property, individual responsibility, and hiring and severance procedures. The Security Policy provides all members of the organization with guidance on all issues of security.

The Security Plan details the implementation of the rules, practices, and procedures that you have deduced are necessary to maintain the desired level of security. In other words, the security plan is the practical application of the security policy.

The Disaster Recovery Plan sometimes is referred to as the contingency plan, business resumption plan, or business continuity plan. Its purpose is to describe, for all persons in the organization, the expected responses to specific security incidents or to other events, such as natural disasters. While you hope that the

Disaster Recovery Plan will never be executed, you must design it so that, when disaster strikes, you **will** be able to minimize its impact on the organization.

Finally, we return to the Security Audit, which can be used to test the effectiveness of the previous steps. To quote Publilius Syrus, writing over 2000 years ago, "It is a bad plan that admits of no modification." The results of the audit will determine the changes that need to be made to policies and plans. These changes then **will** need to be implemented. You also **will** want to conduct tests of a different kind to assess the effectiveness of your Disaster Recovery Plan.

In other words, the process of assessing and planning and implementing is ongoing, as illustrated in Figure 3.1. Managing security is not a one-shot, fix-and-forget proposition. The amount of work involved in the second cycle is considerably reduced due to the groundwork that you will have laid in the first cycle, but the fact remains that it is a lot of work, especially if you do it properly.

Missions and threats

When you start researching the subject of risk analysis, you find that people have gone to great lengths to create a science out of it, one with its own terminology and methodology. It is certainly a field of study that extends far beyond information systems. In plain English, however, risk analysis means knowing what you're up against. In terms

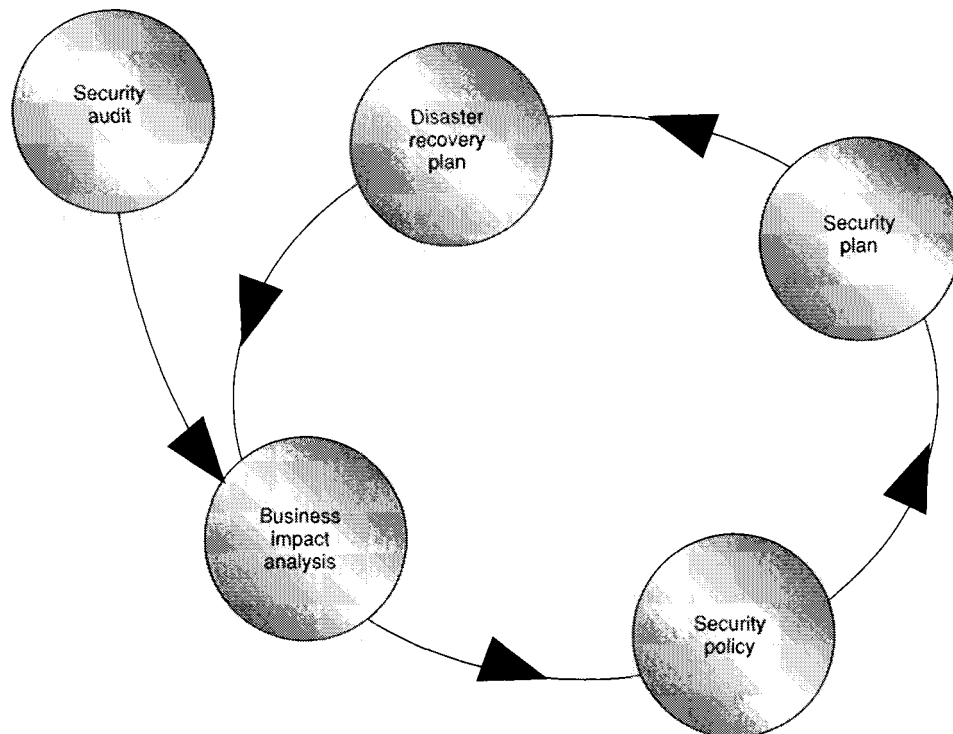


Figure 3.1 The assessment, planning, and implementation cycle.

of personal computer security, it means deciding which threats are the most threatening and which countermeasures are most cost-effective. In appendix A, you will find a comprehensive threat list that might be a useful starting point for this assessment.

One way of categorizing the various approaches to information system risk analysis is to see them as *threat-oriented* or *mission-based*. According to David Snow, writing in *Government Computer News*, threat-oriented risk analysis applies "a pre-defined set of adversary capabilities to determine whether possible system security failures can be exploited." In other words, you devise a list of all the possible attacks, such as the one in appendix A, and assess your system's ability to withstand them.

In contrast, mission-oriented risk analysis methods "attempt to identify all mission hazards early on. The hazards then are used to define the systems' basic security requirements." According to Snow, threat-oriented methods have several drawbacks. First of all, "analysis can only be done when the system is well-defined." This might be appropriate if you are assessing the risks that are posed by a small, existing installation of personal computers, and you could start by checking off a list of possible threats. However, if you are introducing personal computers into an organization or supervising the growth of a personal computer network, a mission-based approach might be more appropriate.

The second disadvantage of threat-oriented thinking is that "only the insecurity of the system is demonstrable — not its security — and there is always uncertainty about additional unidentified impacts on the mission." In other words, any threat list, such as the one in appendix A, is bound to be incomplete.

Mission-oriented risk analysis

Because they offer a framework that you might find useful when you come to assess your organization's security status the following four elements of mission-oriented risk analysis are described: Security Fault Analysis, Threat Analysis, Risk Reduction, and Security Evaluation.

Security Fault Analysis. Security Fault Analysis means looking at the computers, their software, their location, and their purpose in order to identify security loopholes that could lead to a mission hazard. For example, a new personal computer is being installed that will receive inventory reports via modem from other PCs at remote locations. Because the modem will have to be set up to receive data, a path into the computer system is opened up. This will need to be guarded by security measures such as password authentication because unauthorized access to this path could allow someone to steal/damage/compromise data. A weakness in the safeguards placed on the modem connection poses a mission hazard.

Threat Analysis. Threat Analysis means looking at the capabilities of potential adversaries to determine whether or not they could cause or exploit a security failure. Suppose that a modem-equipped personal computer is used to store central inventory records for several branches of a jewelry store. A password is required before an incoming phone connection can be established with the inventory computer. Threat analysis looks at how likely someone is to try to make an unauthorized connection to the computer, what are the chances they will persist in trying to guess the password,

how strong the password is, and what sense or use could be made of the inventory information if the password was broken. A further factor would be an assessment of the impact of corruption or loss of data rather than overt exploitation.

Risk Reduction. If analysis of the weaknesses in the system and the strengths of adversaries reveals risks that are deemed unacceptable, then further risk reduction is necessary. Further countermeasures can be considered, and an assessment made of the relative cost-effectiveness. The chart in Figure 3.2 shows you the basic assumption that the greater the possible loss from a security failure, the greater the justification for further risk reduction.

Security Evaluation. Because absolute security is an unobtainable goal, a given set of security measures can be tested until the weak points are revealed. Evaluation of the security of the system with safeguards in place is an important part of the risk analysis process.

InfoWar and peace of mind

Perhaps the one fact upon which all theories of risk analysis agree is that risk cannot be reduced to zero. Threats cannot be eliminated. They will always exist. However, you can reduce the *likelihood* of occurrence, which is the probability of a loss-causing event. In practical terms, this means that you need to decide what level of risk is acceptable or affordable, realizing that there often will be a trade-off between the two. This is a fact of economics, which was so eloquently described by Lord Butler as "the science which studies human behavior as a relationship between ends and scarce means which have alternative uses." Thus, we might have tighter access control as our "ends" with the department's budget as the "scarce means" and a new office air-conditioning system as just one of the "alternative uses" of those scarce means.




Figure 3.2 Plotting potential loss against risk reduction

Although appendix B lists a wide range of threats, it does not list all of them. Some threats, such as HERF guns and EMP bombs, are not included, although they do exist (see chapter 5 for more details). They generally are considered too "esoteric" to be covered in mainstream computer security texts. However, they are described in depth in *Information Warfare* by American author Winn Schwartau, a book that should be read by everyone with a serious interest in computer security and anyone who depends upon any computer for any information (that's right, I mean everybody). Schwartau goes into depth on a number of subjects that were, until recently, considered science fiction. However, his book is not science fiction. It is a calm and reasonable analysis of the more drastic threats to information that currently exist.



Tempest in a Teacup

When I was writing the first edition of this book, I covered a number of subjects, such as Van Eck Freaking and Tempest technology, that I later found were supposed to be classified. They were subjects about which the general public knew very little, to the point where merely talking about them drew very strange looks. However, they are real, and the threats associated with them have not gone away. Thanks to people like Schwartau, it is possible to read more about them now than ever before.

So why have I omitted from my threat list some of the technology that Schwartau describes? The answer involves four factors: capabilities, intentions, ethics, and economics. When you apply ethics and economics to the consideration of capabilities and intentions, you get different results if you are talking about military decisions as opposed to business decisions. A standard military axiom states that you must defend against your enemies' capabilities, rather than their intentions. In other words, if your enemies have biological weapons, you must have defenses against biological weapons, even if your enemies have repeatedly declared that they have no intention of using them and all the intelligence data suggests that this assertion is genuine.

This thinking implies that arguments such as "it would be madness to use them" and "it would bankrupt them" are null and void. Ethics and economics do not have the same restraining effect in military decisions as they do in a business context. The chart in Figure 3.3 contrasts the relationship between the capabilities of potential adversaries and the acceptability of security expenditures in these two different spheres of human activity.

The point is that a businessperson is strictly confined by economics and cannot act with complete disregard for ethical concerns (unless he or she is prepared to put the organization at risk from costly lawsuits or damaging disclosures). A businessperson must work within the same constraints when assessing threats. Even though it is possible for an individual to build a device with which to severely damage my computer equipment without entering the building where I have my office, I have to question, on ethical and economic grounds, the likelihood of anyone doing this and the value of security measures to protect against it. The questions go something like this: Why would they do this? What would they gain? Could they afford to carry this out?

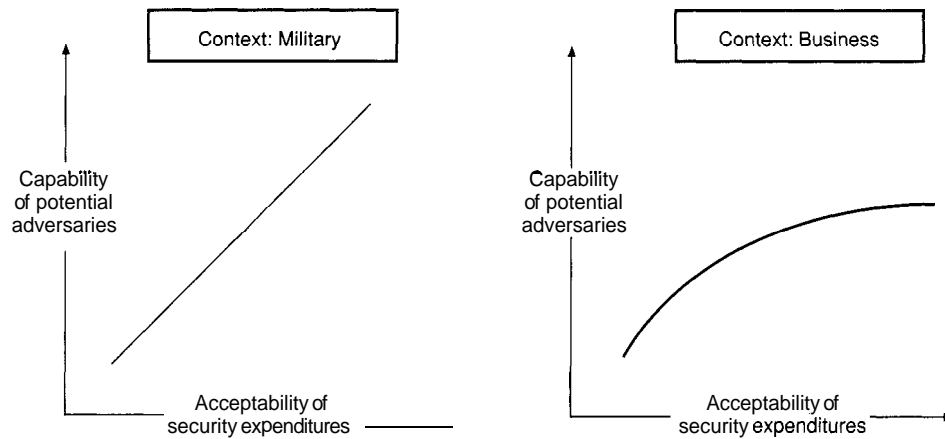


Figure 3.3 Security expenditures relative to the capabilities of potential adversaries.

The bottom line is that I am going to use my scarce means to mitigate the threats that have the highest rates of occurrence. However, I am not going to neglect planning for recovery from disaster, such as fire, flood, or "other," which is the category into which, for the moment, I must place the more esoteric threats. Furthermore, I am going to stay in touch with other security professionals, such as Schwartz, so that I will know as soon as possible when any of these esoteric threats make the transition from possible to probable to everyday.

Basic Risk Analysis

The analysis of risk is an inexact science, dealing as it does with estimated values, chances, and probabilities. A fair amount of literature exists on the theoretical side of this issue. There are even computerized risk analysis systems, such as RiskPAC from CSCI and the Bayesian Decision Support System, both of which are discussed later in this chapter.

Basic questions

Risk analysis is more than figuring the chances of bad things happening. You will need to put a dollar amount on the possible impact of bad things. This then can be used to weigh the cost of defense against the value of what is being defended. Essentially the risk evaluation that you need to perform for personal computer systems should attempt to answer, as reliably as possible, the following questions:

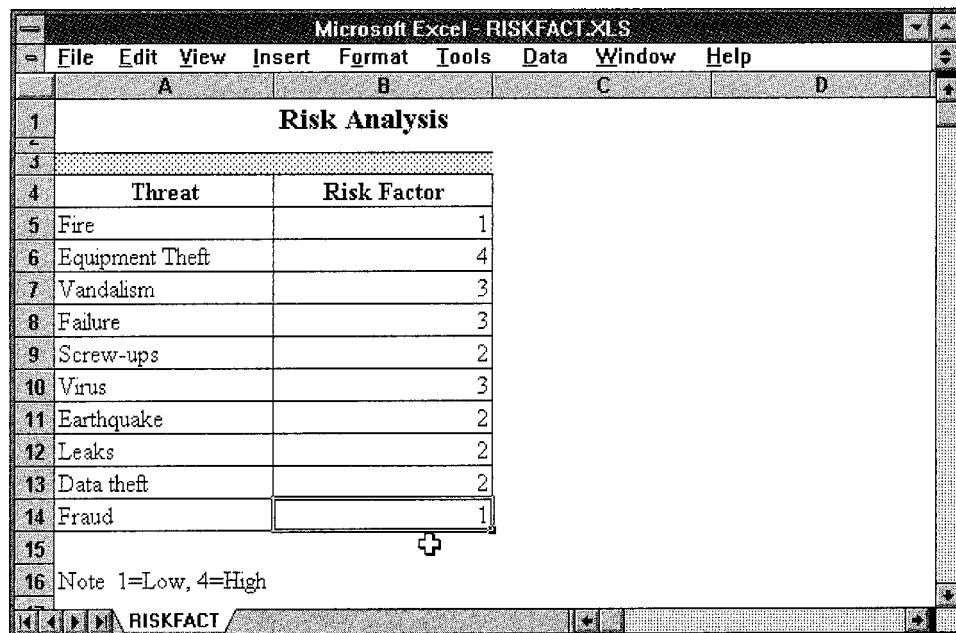
- What are you trying to protect?
- What is its value to you or your organization?
- What are you trying to protect against?
- What is the likelihood of an attack?

A risk analysis example

Risk evaluation involves imagining what could go wrong and then estimating the chances of it actually happening. For each of the possible problems, the question of probability needs to be considered. In this way, the problems and their potential costs can be prioritized and an appropriate plan of action developed. Consider a small consulting company of six people that uses personal computers for its billing and accounting, as well as for the preparation of proposals, contracts, and reports.

The company has few rivals in the area. The company is located on the ground floor of an office building on the fringe of the central business district of a large metropolitan area. The office is not visited by the general public. The company has four PCs, a laser printer, and a modem to call up a commercial online database. The company recognizes that the information that is its lifeblood flows through its personal computer equipment. Thus, the company can define what it is trying to protect as: essential financial data plus core business documents. The value of this information to the company could be rated as "critical."

The office manager sits down with the other employees and fills out the first two columns of a table like the one in Figure 3.4. This is a rather short list, but the company is fairly small and fortunate in that it is not open to the public and does not have extensive connections to outside computers. It answers the question "What are you trying to protect against?" The next question to ask is "What is the likelihood of an attack?" The answer can be described as the *likelihood of occurrence* or the probability of loss-causing event. You can see the company's answers in column three of Figure 3.4.



Risk Analysis	
Threat	Risk Factor
Fire	1
Equipment Theft	4
Vandalism	3
Failure	3
Screw-ups	2
Virus	3
Earthquake	2
Leaks	2
Data theft	2
Fraud	1

Note 1=Low, 4=High

Figure 3.4 A basic risk table

The company is fortunate in having good fire protection in the building and well-trained employees. However, the building location poses problems. The concern over hardware is typical, reflecting what many users see as a relatively low level of hardware reliability in personal computer systems. The lack of concern about user-errors smacks of overconfidence, but here is their summary of risks, ranked by risk factor:

- Common theft: high
- Vandalism: medium
- Equipment failure: medium
- Virus damage: medium
- Screw-ups: low
- Earthquake: low
- Unauthorized access: low
- Data theft: low
- Fire: very low
- Fraud: very low

This ranking will be important when they enter the cost/benefit phase of the evaluation. Note the terms low, medium, and high. While these might seem somewhat informal, they are quite valid in this context.

The next phase for the example company is to evaluate the various threats in terms of potential costs of an occurrence. In other words, they must answer the question "What do we stand to lose?" Here are the results of this evaluation:

Common theft: Equipment is insured, but data is not. Reconstructing data and getting "back to normal" could take up to 10 person-days. Some business might go unbilled and some work might be lost. Potential per incident loss: \$5000.

Vandalism: Same as common theft. Potential per incident loss: \$5000.

Equipment failure: No single piece of equipment would cost over \$3000 to replace. Potential per incident loss: \$3000.

Virus damage: Given current means of early detection, could cause one day disruption for four people. Potential per incident loss: \$1000.

Screw-ups: Given current levels of expertise, could cause half-day disruption for one person. Potential per incident loss: \$500.

Earthquake: Loss of equipment and interruption of business covered by insurance. Potential per incident loss would be limited by deductible: \$5000.

Unauthorized access: Revealing some of our internal comments could cause loss of goodwill. Losing a major client could have a serious impact. Potential per incident loss equal to annual revenue from largest single client: \$50,000.

Data theft: Loss of revenue in one year probably could amount to one quarter of total revenue (currently \$280,000 per annum). Potential per incident loss: \$70,000.

- Fire: Loss of equipment and interruption of business covered by insurance. Potential per incident loss would be limited to deductible: \$5000.
- Fraud: Total cash turnover per month is \$40,000 but this is closely watched. Potential per incident loss: \$10,000.

When you analyze this list, it sheds interesting light on the relative importance of each threat. Consider the spreadsheet in Figure 3.5, which lists threats, risk factors, loss potentials, and loss exposures. You can see that the risk factor column assigns a value to the terms very low, low, medium, and high. The column headed PPIL lists the values identified in the previous list as potential per incident losses, given in units of \$1000. The loss exposure column multiplies the Risk Factors by the Loss Potentials. When you sort this table according to the loss exposure column, you get the very interesting results seen in Figure 3.6.

This tells the company that the highest risk factors do not necessarily create the largest exposures. Another approach would be to replace the high-low risk factors with estimated rates of occurrence. The table in Figure 3.7 shows the estimated incidence of each threat for each year (annualized rate of occurrence or ARO) based on management's best guesses. The results can be termed the annualized loss exposure (ALE), and as you can see, they differ significantly from the estimates based on risk factor. This table will help the company decide which threats should get the most attention.

Microsoft Excel - RISKFACT.XLS				
File Edit View Insert Format Tools Data Window Help				
	A	B	C	D
1	Risk Analysis - Loss Factors			
2				
3				
4	Threat	Risk Factor	PPIL (\$K)	Loss Exposure (\$K)
5	Fire	1	\$5.0	\$5.0
6	Equipment Theft	4	\$5.0	\$20.0
7	Vandalism	3	\$5.0	\$15.0
8	Failure	3	\$3.0	\$9.0
9	Screw-ups	2	\$5.0	\$10.0
10	Virus	3	\$1.0	\$3.0
11	Earthquake	2	\$5.0	\$10.0
12	Leaks	2	\$50.0	\$100.0
13	Data theft	2	\$70.0	\$140.0
14	Fraud	1	\$10.0	\$10.0
15				
16	Note 1=Low, 4=High			

Figure 3.5 A table of threats, risk factors, loss potentials, and loss exposures

Microsoft Excel - RISKFACT.XLS

File Edit View Insert Format Tools Data Window Help

	A	B	C	D
1	Risk Analysis - Loss Factors			
3				
4	Threat	Risk Factor	PPIL (\$K)	Loss Exposure (\$K)
5	Data theft	2	\$70.0	\$140.0
6	Leaks	2	\$50.0	\$100.0
7	Equipment Theft	4	\$5.0	\$20.0
8	Vandalism	3	\$5.0	\$15.0
9	Screw-ups	2	\$5.0	\$10.0
10	Earthquake	2	\$5.0	\$10.0
11	Fraud	1	\$10.0	\$10.0
12	Failure	3	\$3.0	\$9.0
13	Fire	1	\$5.0	\$5.0
14	Virus	3	\$1.0	\$3.0
15				
16	Note 1=Low, 4=High			

RISKFACT

Figure 3.6 The table sorted according to loss exposure

Microsoft Excel - RISKFACT.XLS

File Edit View Insert Format Tools Data Window Help

	A	B	C	D
51	Risk Analysis - ARO			
53				
54	Threat	ARO	PPIL (\$K)	Loss Exposure (\$K)
55	Leaks	0.0140	\$50.0	\$0.700
56	Data theft	0.0070	\$70.0	\$0.490
57	Screw-ups	0.0340	\$5.0	\$0.170
58	Fraud	0.0090	\$10.0	\$0.090
59	Failure	0.0270	\$3.0	\$0.081
60	Equipment Theft	0.0080	\$5.0	\$0.040
61	Fire	0.0075	\$5.0	\$0.038
62	Vandalism	0.0060	\$5.0	\$0.030
63	Earthquake	0.0010	\$5.0	\$0.005
64	Virus	0.0015	\$1.0	\$0.002
65				
66				

RISKFACT

Figure 3.7 Estimated incidence of threats based on management's best guess.

The next step is to place a price tag on potential defensive measures and total up the cost. In order to decide which defensive measure would be applicable, the company first reviewed defenses currently in place:

General: Backup to floppy disks is performed on most machines on an "almost daily" basis with backups stored in locked drawers.

Common theft: Front door lock, blinds closed at night.

Vandalism: Front door lock.

Equipment failure: Tender loving care, regular maintenance, no-smoking rule, "loaner" provisions.

- *Virus damage:* All incoming software vetted on one system, using antivirus software. Public domain and shareware programs only used if from reliable sources.

Screw-ups: Well-trained employees, try to get good temp help when needed.

Earthquake: None apart from fire protection, which is good.

Unauthorized access: Lock on front door. Some of the personal computers have keyboard locks.

Data theft: Lock on front door. Some of the personal computers have keyboard locks.

Fire: Fairly good systems installed at present, no-smoking office.

Fraud: Password protection on accounting system. All reports closely held by accountant.

Next the company came up with suggested defenses that could be added:

General: Fast tape backup system for the PC that runs accounting program, making daily backup less of a chore; get "fast backup" software for other systems and insist on daily backup for all users; get a fireproof filing cabinet for storage of backup media and printed accounting reports, billing slips, and invoices.

Common theft: Install alarm system, arrange office so that computers are not visible from street, attach computers to desks with cable system.

Vandalism: Same as common theft.

Equipment failure: Possibly replace equipment that is getting old. *Virus damage:* Keep anti-virus software up-to-date, watch computer press for virus news.

Screw-ups: Make sure that undo facilities are active in programs used. Install system level rebuild facilities, install unerase programs, install unformat programs, hire temp help from bonded agency specializing in personal computer workers.

Earthquake: fireproof filing cabinet. Same as general.

Unauthorized access: Office alarm system, install system access control software, use passwords on important files, improve handling of keyboard lock keys.

Data theft: Same as unauthorized access.

Fire: The fireproof filing cabinet.

Fraud: Review accounting procedures.

Planning For Risk Reduction					
Additional Defenses	Cost	Description	Cost Shared	ALE	Cost/Benefit
Common theft	\$500	Alarm system	\$250	\$0.040	0.0160
Vandalism	\$0	As above	\$250	\$0.030	0.0120
Equipment failure	\$1,000	Upgrades	\$1,000	\$0.081	0.0081
Virus damage	\$100	Software/magazine	\$100	\$0.002	0.0020
Screw-ups	\$1,500	Tape backup	\$1,500	\$0.080	0.0053
Earthquake	\$250	Fire-proof safe	\$125	\$0.005	0.0040
Fire	\$0	As above	\$125	\$0.038	0.0304
Fraud	\$850	Accounting review	\$850	\$0.090	0.0106
Leaks (unauth/access)	\$100	Software	\$50	\$0.700	1.4000
Data theft	\$0	As above	\$50	\$0.490	0.9800

Notes: Cost Shared column spreads costs of items that have an effect on several areas. The Cost/Benefit column is fairly crude since some items, notably the tape backup unit, offer defense against a wide range of threats.

Figure 3.8 Costs of security measures relative to potential loss/savings

The company is now in a position to weigh the security benefits of each item against the cost, relative to other demands on the budget, bearing in mind that some of the possible measures will be effective on several fronts. In Figure 3.8, you can see a worksheet of costs relative to potential loss/savings.

Beyond basics

From this example, you get an idea of what is involved in a risk evaluation, even if your organization is much larger. The final decision on which measures to adopt will depend on a variety of factors beside cost. It is important to consider the following variables.

Productivity. Personal computers were introduced to increase productivity. Some security measures are so inconvenient that they can impair productivity. Any loss in productivity from a specific measure will need to be weighed against the perceived benefits of the measure.

Feasibility. Some security recommendations that sound good in theory might just not be feasible, given the practical implications. For example, "keep personal computers away from members of the public" is a good idea, but clearly not feasible if the personal computers are being used in a frontline operation like point of sale. Alternatives and compromises will have to be considered. (For example, the security of a personal computer that is exposed to the public might be improved by bolting it down and fitting the keyboard with a protective cover to prevent damage from spills and dirt.)

Aesthetics. All houses would be more secure if they had iron bars across all the windows, but most houses do not. This is because there is a quality-of-life factor to consider in any security decision. There is little point in adopting security measures so ugly that they have a negative impact on morale. Placing trust in employees usually bears more pleasant and profitable fruits than starting out with an adversarial attitude.

Tools and Techniques

The preceding example shows how a small organization can effectively undertake security assessment and risk analysis without any specialized tools, beyond perhaps a spreadsheet for organizing and sorting tables. However, when dealing with larger systems, it is entirely appropriate to bring in outside help, or at least use some specialized tools.

Delphi and expert systems

In fact, the sample company arrived at an *ad hoc* version of a technique known as Delphi, for the oracle in Greek legend. This involves bringing together a group of experts, known as a Delphi panel, who use their combined knowledge and experience to arrive at answers to questions.

Typically the answers are given in numerical form, such as a score or rating, allowing them to be quantified. For example, the panel could be asked to rate the threat of virus infections on a scale from 1 to 10. If there are 10 members on the panel, then a total score out of 100 could be assumed. A series of similar questions about threats could be scored in the same way to create a ranked threat list. However, such a system might not tell the whole story. For example, if all experts rate all threats as very serious the rankings might not be much help in deciding where to apply limited resources. That is where different methods of ranking come into play.

Ranking and quantification

One way of determining the relative seriousness of threats or risks is to use comparison ranking. In other words, each threat can be compared with every other threat, one at a time. Thus, you avoid the problem of grouping all threats as though they were equal. This is the methodology applied by a program called RANK-IT, a limited version of which is provided on the disk that is included with this book. RANK-IT automates the entry of factors that you want to rank, then allows you to fill in the numbers that compare any two factors. You can see this being done in Figure 3.9.

RANK-IT enables you to print out this table in an attractive format or deduce from it a ranked list, which also can be printed out, as shown in Figure 3.10. You will find that this is a very handy tool for a variety of decision-making situations as it requires you to think more carefully about comparative values than a simple rating system.

Further quantifying risks requires formulas, and here are two that I have come across. In the earlier example, you saw that Annualized Loss Exposure = Annualized Rate of Occurrence x Potential Per Incident Loss, or:

$$ALE = ARO \times PPIL$$

14.0	Theft	Theft		Vandalism
		↑ 5	← 4	
13.0	Vandalism	Vandalism		Viruses
		↑ 6	← 3	
9.0	Viruses	↑ 5	← 4	Screw-ups
		↑ 3	← 5	
18.0	Screw-ups	↑ 2	← 7	Screw-ups
		↑ 4	← 6	

F1 -- Help	F6 -- Preview Ranked List	Home -- Top of Sheet
F2 -- Calculate Ranking	F7 -- Clear Ranking Sheet	End -- Bottom of Sheet
F3 -- Preview Ranking Sheet	F8 -- Save/Retrieve/Delete	PgUp -- Cursor right
F4 -- Print Ranking Sheet	F9 -- System Setup	PgDn -- Cursor left
F5 -- Print Ranked List	F10 -- Exit	Tab -- Next column

Figure 3.9 Comparing threats with RANK-IT.

RANKING RESULTS

14.0	Theft	Theft		Vandalism
		↑ 5	← 4	
13.0	Vandalism	Vandalism		Viruses
		↑ 6	← 3	
9.0	Viruses	↑ 5	← 4	Screw-ups
		↑ 3	← 5	
18.0	Screw-ups	↑ 2	← 7	Screw-ups
		↑ 4	← 6	

Figure 3.10 A ranked list of security threats

Another common formula quantifies risk as the product of three factors: threat frequency, the extent to which the organization is exposed to the threat, and the value of the asset at risk. In other words:

$$\text{Risk} = \text{Threat Frequency} \times \text{Exposure Factor} \times \text{Asset Value}$$



Defining Risk Assessment

"Risk assessment is a process that identifies and measures vulnerabilities. It identifies the need for control mechanisms that reduce the threat to organizational functions. It is a systematic means of evaluating the potential for loss given a set of security threats and operational conditions. Its purpose is to provide key personnel with an understanding of what might happen, the severity of the problem, and the need to develop a prevention program. Risk assessment is a key component of an ongoing business continuity program."

Quote from CSCI product literature for RiskPAC.

Assessing value

When it comes to quantifying value there are several factors to consider. The main ones are described in the following sections.

The time factor, Because we all know that time is money and most of us know that computer problems are invariably time-consuming, it should come as no surprise that time is one of the most expensive elements when there is a breach of personal computer security. When you attempt to place a dollar amount on what a security problem might cost, it is important to bear the time factor in mind.

Even if you are insured against most negative aspects of a security breach, time is always involved in filing claims and organizing a return to normal operations. Uninsured losses, such as accidental data loss, can involve extensive recovery time. You might find that some data that is lost is not worth recreating; however, when data has to be recovered, you must put a total cost on the time involved.

You calculate the cost of time lost by multiplying person-hours by dollars-per-hour. Remember that the dollar-per-hour figure is not just an hourly wage but should include employee overhead as well.

The opportunity cost. The term *opportunity cost* has a special meaning in economics; however, in the case of a security problem, it can simply mean what the organization loses because of the problem. This could be a bid deadline or delivery data that is missed, resulting in loss of business. If employees are kept busy repairing a security breach, they will not be able to continue with normal operations, and this can have a snowball effect. One potential benefit of a contingency plan, discussed later in this chapter, is an orderly and speedy return to normal operations after a problem.

The confidence and goodwill cost. Most organizations need both internal and external confidence in order to flourish. The confidence and goodwill of clients and suppliers is essential for continued operation. Employees within an organization need to have confidence that the organization will continue to fulfill its obligations to them. An organization's backers, such as investors and bankers, need to have confidence in its integrity and continued prosperity. Confidence at all these layers can be badly shaken by a serious security breach. The breach does not have to involve massive data loss or serious system damage. The very fact that a breach occurred can be sufficient to threaten the confidence upon which the organization depends.

The cash cost. When you place a dollar amount on the material cost of security breaches, also consider some important accounting principles. Suppose that the IBM PC AT that you bought for \$3000 in 1987 is stolen in 1997. Have you lost \$3000 worth of equipment? In most cases, the answer is no. The replacement cost of a machine with the capabilities of a 1987 IBM PC AT is likely to be less than \$300. Over the years since you bought the old AT, you probably have gained tax benefits from depreciating the original or purchase cost of \$3000. Using straight-line depreciation over 5 years the salvage is \$0 because you will have depreciated the entire cost (5 years x \$600 per year). This probably means that you could not use the loss of the AT as means of reducing your taxes. (In fact, some personal computer equipment can be expensed rather than depreciated, further complicating the tax implications of a loss.)

If you file an insurance claim to replace the AT, you will find out its insured value. This might be based on replacement cost, a depreciated amount, or market value. The last of these is what you could get for the AT on the open market, which probably is only a few hundred dollars. The wide disparity between the different methods of valuation makes it very important to be clear which values you are referring to when you estimate loss exposure, take out insurance, or count your assets. If in doubt, you might want to consult your accountant, particularly if questions of tax are involved.



At Any Cost?

About a year ago, I read a message on CompuServe that sounded too good to be true. Someone was offering to give away brand new full-color Toshiba notebook computers in exchange for used examples of a discounted monochrome laptop model. However, it was a legitimate offer. A medical facility had developed an application that ran on the older models. When they tried it on the new color models, they found that their screens interfered with sensitive medical equipment. In essence, they were locked into a technology that had become priceless through obsolescence. Replacement cost in this case would not follow the normal rules.

Do I need this?

At this point, it should be clear that there is a heck of a lot of work involved in proper security planning. Indeed, you might be wondering if it really is necessary. Do other organizations do this stuff? The answer can be seen in Figure 3.11, which shows some results from the Infosecurity News Annual Industry Survey.

You might wonder how some of them managed to do it, because the same survey reported that 45% of the organizations had annual security budgets of less than \$50,000 despite the fact that 43% estimated that their losses would be more than \$5 million, if computerized data was tampered with, erased, lost, or stolen. In other words, security is chronically and habitually underfunded. Even when it is acknowledged, the cost of data protections is almost always underestimated.

There is a good chance that, when the personal computer systems in your organization were acquired, no provision was made for the cost of performing a full busi-




Figure 3.11 Percentage of organizations that have security and business resumption plans.

ness impact analysis of their use. Like training, security is one of the hidden or ignored costs of computing. Because the personal computer industry has always been price-driven, we are accustomed to reading statements like "Fully loaded Pentium PCs under \$1500" without multiplying that figure by a factor of four to account for licensed software, suitable training, adequate security measures, and administrative overheads such as auditing and disaster recovery planning.

As someone concerned about computer security within your organization, you are likely to find this situation frustrating to say the least. Something that might help you cope with this frustration is mutual commiseration with other security-minded persons. You can get in touch through the various associations listed in appendix E. In chapter 14, I **will** return to the people problems involved in "selling security to management."



The Budget Squeeze

According to Infonetics Research of San Jose, California, "LAN segments will increase 115 percent between now (1993) and 1997, while management budgets are expected to increase by 66%. On the systems side, growth of servers and managed desktops (170% and 160%, respectively) **will** overwhelm the average 40% budget increase."

Computerized risk/impact analysis

It makes sense to apply some of the processing power of the computers that you are trying to protect to the task of analyzing the risks associated with their use. One program that does this is the Bayesian Decision Support System (BDSS), which is described as "a full-support automated package for quantitative risk assessment in the information processing environment." BDSS was several years in the making, using a team of 12 experts with over 125 years of collective experience in information security, risk analysis, and software development. It has been on the market, with regular updates, for more than seven years. According to the people who make BDSS, there are four questions to ask:

What can go wrong?

- How often can it happen?
- What will be the consequences?

How certain are the answers to the first three questions?

Risk analysis means answering these questions. To answer the first, BDSS develops a logical risk model of your information processing environment, based on your responses to the questions it asks. Questions posed by the system are largely determined by answers to previous questions. Using this risk model, BDSS identifies vulnerabilities and maps them to applicable threats.

To answer the second question ("How often can it happen?"), BDSS develops rate of occurrence distributions for various risk scenarios. This is done by integrating three types of information using complex mathematical and logical algorithms:

- The logical risk model
- A large internal database of threat frequencies and safeguard effectiveness

Historical data on the occurrence of various threats at your site

To answer the third question ("What will be the consequences?"), BDSS addresses both tangible and intangible resources, everything from computers to mission/business impact. The user is guided through a process that identifies and assesses the value of resources, as well as the immediate and consequential impact of their loss.

The answer to the fourth question ("How certain are the answers to the first three questions?") is achieved by the Bayesian integration of uncertainty within the BDSS analysis of your logical risk model. This realistic portrayal of the uncertainties in calculated values is a central attribute of the BDSS risk analysis. The BDSS system then goes on to propose safeguards based on its analysis of your risk model. You are free to select the safeguards that you want, and the program then will evaluate the reduction in risk attributable to that selection. The results are graphed, showing the risk levels both before and after the safeguards are applied. This is done with "risk curves" that present the true range of loss exposure and the associated range of probabilities.

After reviewing these "before and after" graphs (without regard to cost at this point), you select the most effective combination of safeguards. Finally, based on cost parameters that you provide, BDSS performs cost-benefit analyses for the safe-

guards that you have identified as being most effective. The risk curves, ranked threats lists, and other summary data help clarify decision issues. The information is presented in a three-part, management-oriented report, with the first part giving an executive summary, the second presenting decision support information, and the third giving technical detail.

To show you what is involved in this particular approach to risk analysis and to give you an idea of the areas that you will need to address in your own analysis, the following paragraphs describe the various BDSS program modules.

Project sizing. This module provides text editing support, enabling the user to draft language for the following areas of the project report:

- Problem statement/background
- Objectives
- Purpose
- Scope and constraints
- Responsibilities
- Risk acceptance criteria

Sample statements are provided in the documentation to be used as models and for reference.

Loss valuation. This module provides input screens in which the user identifies, in detail, all assets within the scope of the assessment. Manufacturer, vendor, model number, quantity, low and high estimates of value as applicable for each tangible item from modems to mainframes and intangibles such as data, business loss, and so on are input to the asset files. Alternatively, you can enter summary-level low and high estimates of loss value for each of the categories of assets.

It is interesting to note that BDSS sees application and system data and software as having two values. First is the replacement cost. Second is the impact on an organization's line of business, mission, or business/mission function. This must be assessed over time if data cannot be processed. The BDSS package helps develop defensible estimates of application and system data and software value ranges by means of proven techniques and rules of thumb. For each category of asset and each threat occurrence frequency distribution, BDSS recognizes associated confidence levels. These levels of confidence assist in representing the uncertainties associated with asset loss values.

Threat vulnerability mapping. This module presents a series of qualitative and quantitative questions. The answers provide BDSS with a map of the vulnerabilities and related threat exposures existing within the scope of the analysis. Each question sequence begins with relatively general questions and proceeds to increasingly detailed questions. The user's answers determine subsequent questions posed by BDSS.

For each threat, BDSS provides the option of entering site-specific threat occurrence frequency data. This data is used to develop Bayesian distributions of threat

frequency. In this module, users can review onscreen or print all of the vulnerabilities organized within their associated threats.

Impact analyzer. This module maps question responses with the question/threat database to develop a risk model of the subject environment, including loss value distributions, threat frequency distributions, and exposure factor distributions, as well as other information pertinent to risk quantification. This module has no user interface. It is internally executed as necessary. Output, which is provided to the evaluate and revise module and the risk analyzer, includes six items for each of the affected asset categories of each threat: derived annualized rate of occurrence (ARO), national ARO frequency distributions, summary single loss exposure (SLE) distributions, summary exposure factor distributions, derived SLE distributions, and exposure factor distributions.



Good Advice

"It is less costly to prevent a disaster than to recover from one. If you can identify the vulnerabilities in your organization's physical and information security and eliminate the vulnerabilities that can be eliminated in a cost-effective manner, you will reduce the probability of a major disruption to your organization."

Quote from CSCI product literature for RiskPAC.

Risk analyzer. This module accepts the risk model produced by the impact analyzer and submits the data to a series of sophisticated statistical algorithms, including Bayesian algorithms addressing uncertainty. This module has no user interface and is internally executed as necessary. Output from this module is a file of data that represents the risk curve and the average annualized loss exposure (ALE) for each threat.

Safeguard analyzer. This module affects the selection and application of safeguards to vulnerabilities derived in the threat vulnerability mapping module. When you enter this module, it analyzes your responses to questions in the threat vulnerability mapping module, presents indicated vulnerabilities, and suggests appropriate safeguards. Users can select the safeguards that they want to consider implementing. Loss exposures associated with the affected threats are reduced accordingly. Results of safeguard analysis can be previewed in the evaluate and revise module by looking at a single threat with all suggested safeguards applied or, conversely, at a single safeguard with all affected threats represented.

Safeguard cost/benefit. In this module, the costs of selected safeguards are developed. The cost/benefit analysis is completed by BDSS through present value techniques applied to the projected implementation and maintenance costs of subject safeguards versus expected savings (loss reduction) over the life of the safeguards. Then, based on expected loss reduction, users select those safeguards that best address the vulnerabilities and associated loss exposures. Alternatively, users can allow BDSS to present all analyzed safeguards for management to review and select for implementation.

Evaluate and revise. This module allows evaluation of the tabular frequency distribution and exposure factor data from the impact analyzer. It also permits review of the risk curves for any given threat, combination of threats, or safeguards. The nonreduced risk curve for any given threat or combination of threats can be viewed. If the safeguard analyzer has been executed, users can review the loss reduction achieved by applying safeguards via the superimposed reduced risk curve. This module also can provide a list of the threats ranked according to their loss potential (contribution to overall risk as denoted by the average ALE for each threat) before and after the application of safeguards.

Report generator. This module assembles all of the information provided by the users, as well as calculated information, and generates a set of reports. The reports are described here because they give you a good idea of what type of information needs to be assembled to make fully informed security planning decisions. The executive summary reports present:

- Project problem statement and objectives.

- Scope and constraints.

- Overview of the Bayesian Decision Support System methodology.

- Discussion of key vulnerabilities and their significance

- Graphic display depicting the unacceptable risk region, a summary nonreduced risk curve for all threats with a superimposed summary reduced risk curve that represents all threats with risk reduction measures applied, plus an incremental curve representing summarized safeguard costs.

- General recommendations.

The executive decision support report provides the decision-making executive (and other interested parties) with all the information needed to make well-informed and defensible decisions regarding the purchase and/or development and implementation of safeguard measures. The parts of this report are:

Introduction: Contains the problem statement, objectives, scope and constraints, and recommendations rationale as entered originally in the project sizing module.

Approach: Describes the BDSS methodology and its application

Graphic threats *summary*: Summarizes the safeguards that are applied to summarized threats, and the resulting reduced risk curve and safeguard costs curve are depicted superimposed on the summarized threats' nonreduced risk curve.

Threat graphics: Graphs of the nonreduced risk curve for each single threat with superimposed reduced risk curve and associated costs curve that represents applicable recommended safeguards, with supporting tabular values for the subject threat and mitigating safeguard.

- Safeguard graphics: Charts each recommended safeguard via a reduced summary risk curve for all affected threats, superimposed on the nonreduced summary risk curve for the same threats, with supporting tabular values and costs for the subject safeguard and affected threats.

Asset inventory summary: Lists the assets, summarized by category

Ranked threats summaries: Lists all of the threats identified in the assessment and their associated loss exposures before safeguards are applied, and the same threat list resequenced as appropriate after safeguards are applied to reflect their reduced loss exposure.

Detailed recommendations: User-generated narrative detailing the safeguard recommendations as desired.

The technical analysis reports provide all the information, fully detailed, that was generated by and supported the Bayesian Decision Support System:

Detailed asset inventory: Organized by category.

Threat/vulnerabilities summary: A brief discussion of each vulnerability and associated threats identified in the assessment before and after safeguards are applied.

Threats and loss analysis: A list of all threats and identification of those threats having the potential for impacting the subject information systems environment, based on responses during the threat vulnerability mapping. Both national and localized ARO distributions are presented for each threat. The assets potentially impacted by each threat also are identified here along with the associated loss values and exposure factor distributions.

Safeguards selection and *cost/benefit* analysis: Correlates vulnerabilities with applicable threats and presents an analysis of that information with the selected combination of safeguards that will minimize the loss potential for the subject information systems environment.

Securing your analysis. The information assembled in a typical risk assessment effectively presents a map of vulnerabilities that a malcontent or other individual with destructive intentions could use to damage or destroy the organization. A programmed analysis such as BDSS must allow the owner or project manager to control access to the various projects and "what ifs" within a project. With its system functions module, BDSS provides this function, allowing the owner to prevent access to specific information by all except those with a need to know.

RiskPAC

By now it should be clear that a thorough analysis of security risks is time-consuming. The right software can help reduce the amount of time required and BDSS is not the only package available. Another "interactive, automated risk assessment program" is RiskPAC from Computer Security Consultants, Inc. (CSCI). This is an expert system that incorporates the extensive consulting experience of CSCI into a questionnaire that forms the basis of the user interface.

The company's expertise in information security, computer technology, and auditing are captured in a series of specific questionnaires. You can use these as is or tailor them to suit your environment. As you answer the questions, rules developed by CSCI score the answers and provide a set of descriptive conclusions and recommen-

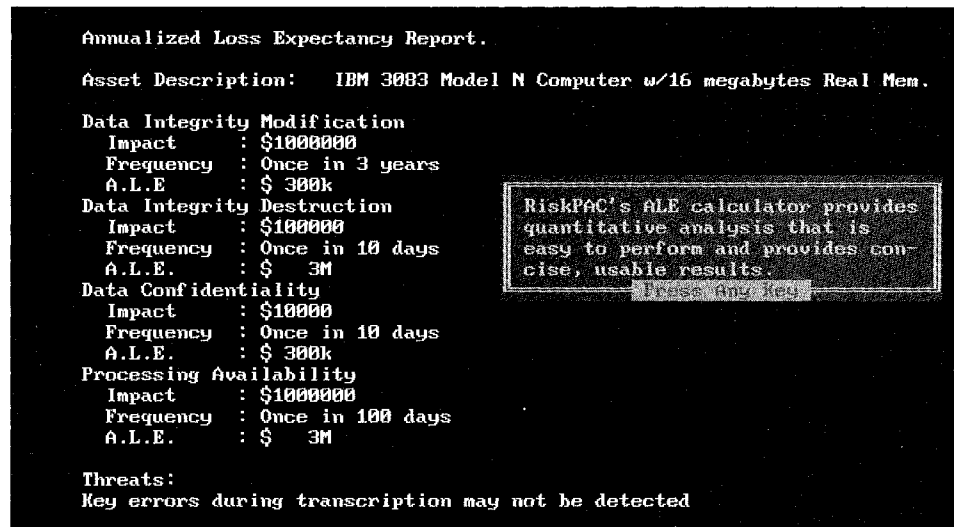


Figure 3.12 The RiskPAC Annualized Loss Expectancy Calculator.

dations for action. If you have your own inhouse experts, you can use RiskPAC System Manager to create your own custom questionnaires that focus on your corporate philosophy, organizational standards, or semantics. By letting your PC do much of the work involved in organizing and collating information, you can take a lot of the drudgery out of risk analysis. For example, in Figure 3.12, you can see the RiskPAC Annualized Loss Expectancy Calculator at work.

A companion CSCI product is RecoveryPAC, which helps disaster-recovery plan developers to produce effective, tailored, testable plans. It integrates the power of a relational database with graphics-oriented project-planning programs to provide a solid set of tools for collecting and updating data, defining and sequencing tasks, and tracking and evaluating tests in real time. Over 100 report formats ensure that recovery plans are easily interpreted and distributed throughout the organization. Reports can be produced on screen, on paper, or in the ASCII file format accessible through most word processors. RecoveryPAC is targeted toward medium to large data centers or total organizations. A downsized version called RecoveryPAC II also is available. It is targeted toward small data centers or single organizational departments. This version lacks project management or testing modules and offers fewer report formats but includes all other RecoveryPAC capabilities.

Security Policy

Like a personnel or employment policy, a security policy is an essential document for a company that relies on personal computers. A *security policy* can be defined as a set of rules, principles, and procedures that regulate how an organization manages, protects, and controls personal computer resources and the information that they contain. By assembling such a policy document, an organization is forced to face up

to questions of security. Obviously, a security policy does not provide security in itself. However, it is a very effective tool, a measure of current levels of security, a standard that can measure compliance, and basis for further improvements.

Policy needs, types, and strategies

Executives and managers are faced with many choices in directing the protection of computer assets. Some choices can be based upon quantifiable trade-offs, but others involve competing trade-offs, questions of organizational strategic direction, or other parameters that do not lend themselves to quantitative analysis. In making these choices, policy is established for an organization and then is used as the basis for protecting resources, both information and technology, and guiding employee behavior.



Your Money at Work

Like risk/business impact analysis, security policy making can be a daunting task. Fortunately, there is a lot of help available, some of it free. The U.S. government has had to deal with issues of computer security for longer than most commercial organizations. Various agencies have spent a lot of time and effort developing security policies and, because this work is funded by us taxpayers, it is available to us in the form of public domain documents. Appendix E provides a list of sources for this literature, some of which is well worth obtaining as it applies equally to many nongovernmental organizations. Indeed, the next few sections are based on a CSL Bulletin from NIST, so you can get some idea of the quality of these documents.

Familiarity with various types of policy will aid managers in addressing computer security issues important to the organization. Effective policies ultimately result in the development and implementation of a better computer security program and better protection of systems and information. Thus, I have described four types of computer security policy, their components, and aspects of policy implementation. Comparison of your organization's computer security policies to the types described here will assist you in determining if your policies are comprehensive and appropriate. Briefly, the four types are: program-level policy, program-framework policy, issue-specific policies, and system-specific policies.

Program-level policy is used to create an organization's computer security program. Organizations need program-level policy to establish the security program, assign program management responsibilities, state organization-wide computer security purpose and objectives, and provide a basis for compliance. Program-level policy typically is issued by the head of the organization or another senior official, such as the top management officer.

Program-framework policy establishes the organization's overall approach to computer security—in other words, its computer security framework. Program-framework policies provide organization-wide direction on broad areas of program implementation. For example, they might be issued to assure that all components of

an organization address contingency planning or risk analysis. They are appropriate when an organization can yield benefits from a consistent approach. Typically, program-framework policies are issued by a manager with sufficient authority to direct all organization components on computer security issues. This might be the organization's management official or the head of the computer security program.

Issue-specific policies address specific issues of concern to the organization. Issue-specific policies identify and define specific areas of concern and state the organization's position. Depending upon the issue and attendant controversy, as well as potential impact, issue-specific policy might come from the head of the organization, the top management official, the Chief Information Officer, or the computer security program manager.

System-specific policies focus on policy issues that management has decided for a specific system. System-specific policies state the security objectives of a specific system, define how the system should be operated to achieve the security objectives, and specify how the protections and features of the technology will be used to support or enforce the security objectives. A system refers to the entire collection of processes, both automated and manual. System-specific policy normally is issued by the manager or owner of the system (which could be a network or application) but might originate from a high official, particularly if all impacted organizational elements do not agree with the new policy.

Program-level policy

Program-level policy establishes the computer security program and its basic framework. This high-level policy defines the purpose of the program and its scope within the organization, assigns responsibilities for direct program implementation (to the computer security organization) as well as responsibilities to related offices (such as human resources), and addresses compliance issues. The components of program-level policy should include: purpose, scope, responsibilities, and compliance.

Purpose. A clear statement of the purpose of the program includes defining the goals of the computer security program as well as its management structure. Security-related needs—such as integrity, availability, and confidentiality—can form the basis of organizational goals established in policy. For instance, in an organization responsible for maintaining large mission-critical databases, reduction in errors, data loss, or data corruption might be specifically stressed. In an organization responsible for maintaining confidential personal data, however, goals might emphasize stronger protection against unauthorized disclosure.

The program management structure should be organized to best address the goals of the program and respond to the particular operating and risk environment of the organization. Important issues for the structure of the central computer security program include management and coordination of security-related resources, interaction with diverse communities, and the ability to relay issues of concern to upper management. The policy also could establish operational security offices for major systems, particularly those at high risk or most critical to organizational operations.



Raise Policy Profile

To be effective, policy requires visibility. Visibility aids implementation of policy by helping to ensure that knowledge of the policy is diffused throughout the organization. Use management presentations, videos, panel discussions, guest speakers, question/answer forums, and newsletters, as resources permit.

Scope. The scope specifies which resources (including facilities, hardware, and software), information, and personnel the program covers. In many cases, the program will cover all systems and personnel, but this is not always true. In some instances, a policy might name specific assets, such as major sites and large systems. Often tough management decisions arise when defining the scope of a program, such as determining the extent to which the program applies to contractors and outside organizations utilizing or connected to the organization's systems. (*Note:* The Computer Security Act of 1987 requires federal agencies to address the security of all federal interest systems).

Responsibilities. This component addresses the responsibilities of officers and offices throughout the organization, including the role of line managers, applications owners, users, and the data-processing organization. The policy statement should distinguish between the responsibilities of computer services providers and the managers of applications utilizing the computer services. It also can serve as the basis for establishing employee accountability. Overall, the program-level assignment of responsibilities should cover those activities and personnel who will be integral to the implementation and continuity of the computer security policy.

Compliance. This component authorizes the use of specified penalties and disciplinary actions for individuals who fail to comply with the organization's computer security policies. Because the security policy is a high-level document, penalties for various infractions normally are not detailed here. However, the policy might authorize the creation of compliance structures that include violations and specific penalties. Infractions and associated penalties usually are defined in issue-specific and system-specific policies. When establishing compliance structures, consider that violations of policy can be unintentional on the part of employees. For example, non-conformance can be due to a lack of knowledge or training.

Program-framework policy

Program-framework policy defines the organization's security program elements that form the framework for the computer security program and reflect decisions about priorities for protection, resource allocation, and assignment of responsibilities. Criteria for the types of areas to be addressed as computer security program elements include, but are not limited to:

- Areas for which there is an advantage to the organization by having the issue addressed in a common manner
- Areas that need to be addressed for the entire organization

Areas for which organization-wide oversight is necessary

Areas that, through organization-wide implementation, can yield significant economies of scale.

The types of areas addressed by program-framework policy vary within each organization as does the way in which the policy is expressed. Some organizations issue policy directives, while others issue handbooks that combine policy, regulations, standards, and guidance. Many organizations issue policy on "key" areas of computer security, such as life-cycle management, contingency planning, and network security.

Keep in mind the criteria stated earlier for the types of areas that should be addressed in program-framework policy. If the policy (and its implementing standards and guidance) is too rigid, cost-effective implementations and innovation could be stifled. As an example of program-framework policy, consider a typical organization policy on contingency planning. The organization might require that all contingency plans categorize criticality of processing according to a standard scale. This will assist the organization in the preparation of a master plan (for use if the organization's physical plant is destroyed) by facilitating prioritization across intra-organizational boundaries.

Policy in these areas normally applies throughout the organization and usually is independent of technology and the system or application. Program-framework policies might be comprised of components similar to those contained in program-level policy but also might be in a very different format (for example, in organizational handbook directives).

Issue-specific policy

Issue-specific policies focus on areas of current relevance and concern, perhaps even controversy. Program-level policy usually is broad enough that it requires little modification over time. Conversely, issue-specific policies require more frequent revision due to changes in technology and related factors. As new technologies develop, some issues diminish in importance while new ones continually appear. It might be appropriate, for example, to issue a policy on the proper use of a cutting-edge technology, such as World Wide Web pages on the Internet, the security vulnerabilities of which still are largely unknown.

A useful structure for issue-specific policy is to break the policy into its basic components: statement of an issue, statement of the organization's position, applicability, roles and responsibilities, compliance, and points of contact. Other topic areas can be added as needed.

Issue statement. The issue statement defines the issue, with any relevant terms, distinctions, and conditions. For example, an organization might want to develop an issue-specific policy on the use of "unapproved software," which might be defined to mean any software not approved, purchased, screened, managed, and owned by the organization. Additionally, applicable distinctions and conditions might possibly be included, for instance, software privately owned by employees but approved for use at work and for software owned and used by other businesses under contract to the organization.

Statement of the organization's position. The statement of the organization's position states the organization's position on the issue. To continue the example of unapproved software, the policy would state whether use of unapproved software is prohibited in all or some cases, whether or not there are further guidelines for approval and use, or whether case-by-case exceptions will be granted, by whom, and on what basis.

Applicability. This describes where, how, when, to whom, and to what a particular policy applies. For example, the hypothetical policy on unapproved software might apply only to the organization's own onsite resources and employees and not to contractor organizations with offices at other locations. Additionally, the policy's applicability to employees traveling among different sites or working at home who will transport and use disks at multiple sites might require clarification.

Roles and responsibilities. This assigns roles and responsibilities within the organization for specific policies. To continue the software example, if the policy permits unapproved software privately owned by employees to be used at work with appropriate approvals, then the approving authority would need to be identified. An office responsible for compliance also could be named.

Compliance. This gives descriptions of the infractions that are unacceptable and states the corresponding penalties. Penalties must be consistent with organizational personnel policies and practices and need to be coordinated with appropriate officials, offices and, perhaps, employee bargaining units.

3

Coordination Is Key

Computer security policy should be integrated into and consistent with other organizational policies, such as personnel policies. One way to help ensure this is to thoroughly coordinate policies during development with other offices in the organization.

Points of contact and supplementary information. This provides the name of the appropriate individuals to contact for further information and lists any applicable standards or guidelines. For some issues, the point of contact might be a line manager; for other issues, it might be a facility manager, technical support person, or system administrator. For yet other issues, the point of contact might be a security program representative. Using the software example, employees need to know whether the point of contact for questions and procedural information would be the immediate superior, a system administrator, or a computer security official.

System-specific policy

Program-level policy and issue-specific policy both address policy from a broad level, usually encompassing the entire organization. System-specific policy, on the other hand, is much more focused, because it addresses only one system. Many security

policy decisions apply only at the system level. Some examples include:

Who is allowed to read or modify data in the system? Under what conditions can data be read or modified?

- Are users allowed to dial into the computer system from home or while on travel?

To develop a comprehensive set of system security policies, use a management process that derives security rules from security goals. Consider a three-level model for system security policy: security objectives, operational security, and policy implementation.

Security objectives. First, define the security objectives. While this process might start with an analysis of the need for integrity, availability, and confidentiality, it cannot stop there. A security objective must be more specific, concrete, and well-defined. It also should be stated so that it is clear that the objective is achievable. The security objectives should consist of a series of statements that describe meaningful actions about specific resources. These objectives should be based on system functional or mission requirements but should state the security actions that support the requirements.

Operational security. Next lay out the operational policy that gives the rules for operating a system. Following the same integrity example, the operational policy would define authorized and unauthorized modification: who (by job category, by organization placement, or by name) can do what (modify, delete, and so on) to which pieces of data (specific fields or records) and under what conditions. Managers need to make decisions in developing this policy because it is unlikely that all security objectives will be fully met. Cost, operational, technical, and other constraints will intervene.

Consider the degree of granularity needed for operational security policies. Granularity refers to how specific the policy is with regard to resources or rules. The more granular the policies, the easier to enforce and to detect violations. A policy violation might indicate a security problem. In addition, the more granular the policy, the easier to automate policy enforcement.

Consider the degree of formality that you want in documenting the policy. Once again, the more formal the documentation, the easier to enforce and to follow policy. Formal policy is published as a distinct policy document; less formal policy might be written in memos. Informal policy might not be written at all. Unwritten policy is extremely difficult to follow or enforce.

On the other hand, very granular and formal policy at the system level also can be an administrative burden. In general, good practice suggests a granular formal statement of the access privileges for a system due to its complexity and importance. Documenting access controls policy makes it substantially easier to follow and to enforce. Another area that normally requires a granular and formal statement is the assignment of security responsibilities.

Some less formal policy decisions might be recorded in other types of computer security documents such as risk analyses, accreditation statements, or procedural manuals. However, any controversial, atypical, or uncommon policies might need

formal policy statements. Atypical policies would include any areas where the system policy is different from organization policy or from normal practice within the organization, either more or less stringent. They also should contain a statement explaining the reason for deviation from the organization's standard policy.

Policy implementation. Determine the role technology will play in enforcing or supporting the policy. Security normally is enforced through a combination of technical and traditional management methods. While technical means are likely to include the use of access control technology, there are other automated means of enforcing or supporting security policy. For example, technology can be used to block telephone systems users from calling certain numbers. Intrusion detection software can alert system administrators to suspicious activity or take action to stop the activity. Personal computers can be configured to prevent booting from a floppy disk.

Automated security enforcement has advantages and disadvantages. A computer system — properly designed, programmed, and installed — consistently enforces policy, although no computer can force users to follow all procedures. In addition, deviations from the policy sometimes might be necessary and appropriate. This situation occurs frequently if the security policy is too rigid



Avoid Policy Overload

Introduce computer security policies in a manner that ensures that management's unqualified support is clear, especially in environments where employees feel inundated with policies, directives, guidelines, and procedures. The organization's policy is the vehicle for emphasizing management's commitment to computer security and making clear their expectations for employee performance, behavior, and accountability.

Policy conclusions

Formulating viable computer security policies is a challenge for any organization and requires communication and understanding of the organizational goals and potential benefits to be derived from policies. Through a carefully structured approach to policy development, which includes the delegation of program management responsibility and an understanding of program-level, program-framework, issue-specific, and system-specific policy components, your organization can achieve a coherent set of policies. These will help produce a framework for a successful computer security program.



Another Plug for Policy

Everyone acknowledges that one of the main reasons for the persistence of computer crime is the failure to successfully prosecute offenders. One of the main reasons for this failure is a lack of clearly defined policy. The problem was neatly summarized by Don Ingram, a computer-savvy district attorney in Alameda, California, when he said this about inappropriate computer behavior: "If you can't define it, how are you going to prosecute it."

Getting to Grips with Policy

In all but the smallest organizations, policy is written at a broad level. For this reason, organizations also develop standards, guidelines, and procedures that offer users, managers, and others a clearer approach to implementing policy and meeting organizational goals. Standards and guidelines specify the technologies and methodologies to be used to secure systems. Procedures are yet more detailed steps to be followed to accomplish particular security-related tasks. Standards, guidelines, and procedures can be disseminated throughout an organization via handbooks, regulations, or manuals.

Standards

Organizational standards specify uniform use of specific technologies, parameters, or procedures when such uniform use will benefit an organization. Standardization of organization-wide identification badges is a typical example, providing ease of employee mobility and automation of entry/exit systems. Standards normally are compulsory within an organization.



Policies Pay Off

Clearly defined security policies make sense for many reasons. In 1991, a California prosecutor observed: "We probably turn down more cases (involving computer break-ins) than we charge, because computer system proprietors haven't made clear what is allowed and what is not."

Guidelines

Guidelines assist users, systems personnel, and others in effectively securing their systems. The nature of guidelines, however, immediately recognizes that systems vary considerably and imposition of standards is not always achievable, appropriate, or cost-effective. For example, an organization guideline might be used to help develop system-specific standard procedures. Guidelines often are used to help ensure that specific security measures are not overlooked, although they can be implemented, and correctly so, in more than one way.

Procedures

Procedures normally assist in complying with applicable security policies, standards, and guidelines. They are detailed steps to be followed by users, system operations personnel, or others to accomplish a particular task (for example, preparing new user accounts and assigning the appropriate privileges). Some organizations issue overall computer security "manuals," "regulations," "handbooks," or similar documents. These might mix policy, guidelines, standards, and procedures because they are closely linked. While manuals and regulations can serve as important tools, they are most useful when they clearly distinguish between policy and its implementation (sometimes a difficult process). This promotes flexibility and cost-effectiveness by offering alternative implementation approaches to achieving policy goals.

Further assistance

If all of this planning and policy making sounds daunting, bear in mind that there are numerous sources of assistance. For a start, appendix C provides examples of security policy documents, and appendix F points to further resources. Remember that, although your organization is unique in some ways, its computer security needs are similar to those of many other organizations. This means that many of the policies that you need to write already have been written.

One of the best commercial sources is a package called *Information Security Policies Made Easy* by Charles Cresson Wood. This is not a software program but a "cookbook" that contains 600 predefined policies. From the lists in Figures 3.13 and 3.14, you can see that just about every imaginable aspect of PC and LAN security has been addressed.

These are not policy statements that you have to painstakingly retype. They also are supplied on disk, for either Mac or PC, enabling you to copy-and-paste only what you need, then edit it to your exact specifications. Although this package costs considerably more than your average book, its comprehensive treatment of the subject, together with the ability to directly feed the text into your own documents, makes it a very worthwhile investment.

Disaster-Recovery Planning

Even if everyone adheres to well-defined security procedures, problems can arise. Natural disasters are notably difficult to predict or prevent. In the disaster-recovery plan (DRP), you lay down procedures to be followed in the event of a serious problem. A comprehensive treatment of disaster-recovery planning is beyond the scope of this book. Two books worth reading on this subject are *Disaster Recovery for LANs* by Regis J. Bates (McGraw-Hill, 1994) and *Enterprise Disaster Recovery Planning* by Michael Miora (McGraw-Hill, 1996).

Basic DRP

You might begin consideration of a DRP by looking over a list of all possible problems taken from your risk/business impact analysis. The next step is to decide who should do what when problems materialize. In larger organizations, there are two aspects to the DRP: operational and administrative. At the operational level, each user should know what to do when a problem arises. At a minimum, this knowledge will include an answer to the classic question, "Who you gonna call?" At the administrative level, the plan will cover such questions as where to hire replacement equipment, how to restore data from backups, what documentation is needed for an insurance claim, and how to recreate data in the event both files and backups are destroyed.

Advanced DRP

When you are dealing with large-scale, multiserver networks that are running mission-critical applications, the DRP becomes even more complex and of vital importance to the organization. You can see this from the flowchart in Figure 3.15, which was

Personal Computer Security Policy Test

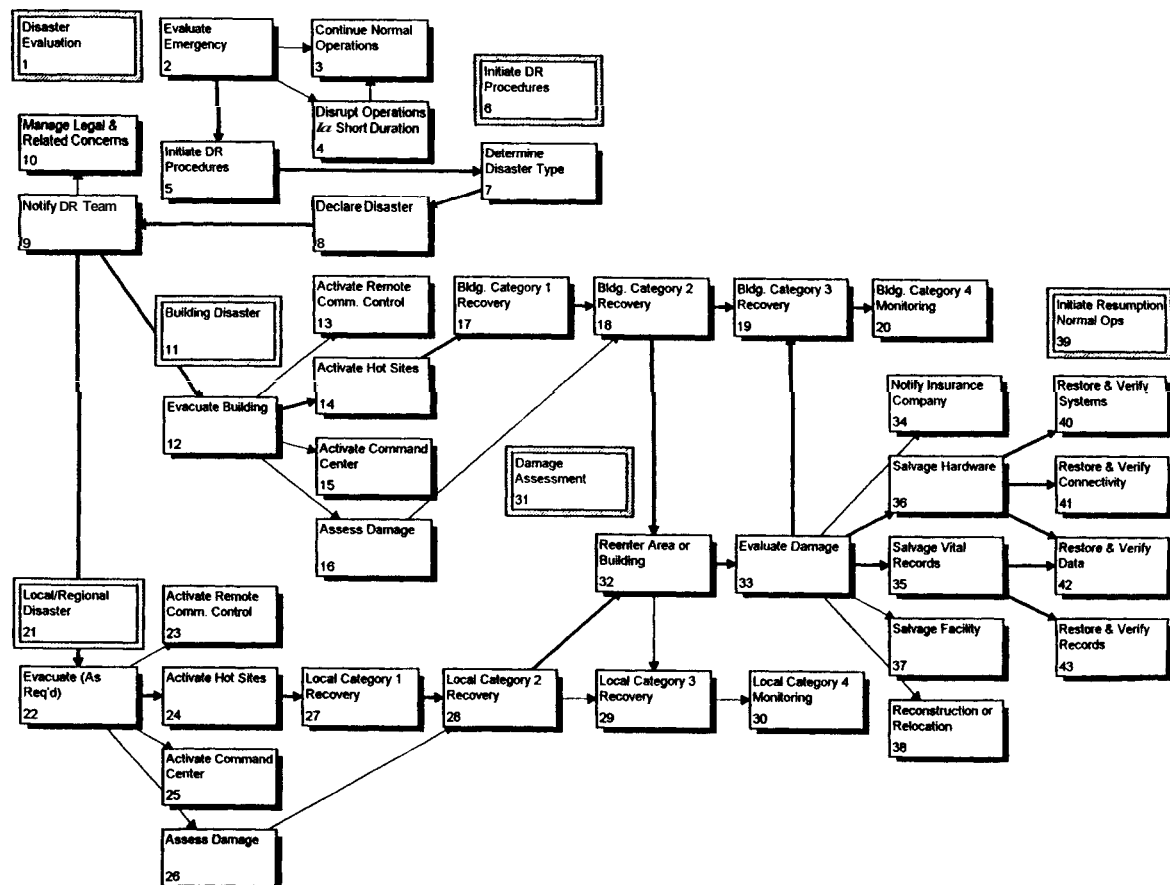
19. Storage of Passwords in Readable Form
36. User-ID and Password Required for Computer-Connected Network Access
43. Automatic Log-Off Process
44. Leaving Sensitive Systems Without Logging Off
45. Logging-Off Personal Computers Connected to Networks
47. Games May Not Be Stored or Used on Company X Systems
82. Inclusion of Security Relevant Events in System Logs
103. Initial Back-Up Copies of Microcomputer Software
113. Use of Higher Level Programming Language
117. Documentation Required for All Production Business Systems
131. Loading External Programs Onto Network-Connected Computers
132. Control Over Movement of Software From Development to Production
139. Periodic Review of Production Operating System Changes
142. Special Approval Required for Production Software Package Changes
154. Smoking, Eating, and Drinking in the Computer Machine Room
162. Workers May Make Multiple Copies Only if Reasonable and Customary
163. Periodic Review of Software Licensing Agreements
165. When Making Additional Copies of Software is Permissible
166. Tools Used to Break Systems Security Prohibited
202. Privacy of Personal Files Stored on Computers and in Desks
284. Clean Desks and Working Areas
286. Storage of Sensitive Information on Personal Computers
287. Storage of Sensitive Information When Not in Use
293. Use of Hard-Disk Drives for Storage of Sensitive Information
294. Commingling of Sensitive and Non-Sensitive Information
322. Browsing on Company X Systems and Networks Prohibited
323. Power-Down Required for Systems Processing Sensitive Information
325. Positioning of Computer Display Screens with Respect to Windows
334. Computing Environment Supporting Equipment Required
335. Power Conditioning Equipment Required for All Microcomputers
336. Static Electricity Protection Equipment and Local Conditions
340. Compliance with Standards Required for Emergency/Disaster Support
352. Annual Inventory of Information Systems Hardware, Software, Etc.
358. What Data to Back-Up and Minimum Back-Up Frequency
360. Two Copies of Sensitive, Critical, or Valuable Information
361. Management Review of End-User Back-Up Process
363. Automatic Back-Up to Local Area Network Server
379. Source and Data Labels Required for Major Decision Input
388. Review Required for Important Computer Analysis Done by Individuals
389. Proper Controls for Data Used to Arrive at Decisions Involving \$10,000
412. Prior Approval Required for All Communications Line Changes
413. Prior Approval Required for Set-Up of Multi-User Systems
416. Use of Computer Equipment Belonging to Workers on Company Property
423. Dial-Up Connections Must Always Utilize Firewalls
428. Secret Information Must be Encrypted When Not in Archive Use
429. Data Stored on Hard Disk Drives Must be Encrypted
451. Modems on Workstations Connected to Internal Networks
455. Down-Loading Sensitive Information Prohibited Without Permission
549. Each Department Must Have an Information Security Liaison
567. Physical Access Control for Areas Containing Sensitive Information
576. Physical Security or Encryption Required for All Sensitive Information
577. Property Pass for Removal of All Computer and Communications Gear
579. Provision of Lockable Metal Furniture to Staff Working at Home
588. Physical Security Measures for Computers and Communications Systems
601. Computer-Assisted Equipment Tracking
602. Positioning Workstations to Reduce Risk of Overlooking

Figure 3.13 Personal Computer Security Policy List (*Reproduced with the permission of Charles Cresson Wood*)

Local Area Network (LAN) Security Policy List

1. Minimum Password Length
2. Cyclical Passwords Prohibited
16. Limit on Consecutive Unsuccessful Attempts to Enter a Password
17. Single Sign-On Process
32. Password Sharing Prohibition
33. Forced Change of All Passwords
36. User-ID and Password Required for Computer-connected Network Access
37. Unique User-ID and Password Required
39. Security Notice in System Log-In Banner
40. Disclosure of Information in System Log-In Banner
42. Prohibition of Multiple Simultaneous On-Line Sessions
43. Automatic Log-Off Process
45. Logging-Off Personal Computers Connected to Networks
48. Prohibition Against Non-Approved System Users
50. Granting User-IDs to Outsiders
59. Re-Use of Unique User-IDs Prohibited
61. Restriction of Special System Privileges
63. Default User Privileges and Need for Explicit Approvals
64. Restriction of Third Party Dial-Up Privileges
67. Default to Denial of Access Control Privileges
68. End-User Access to Operating System Commands
82. Inclusion of Security Relevant Events in System Logs
84. Accountability and Traceability for All Privileged Systems Commands
92. Clock Synchronization for Accurate Logging of Events on Network
117. Documentation Required for All Production Business Systems
163. Periodic Review of Software Licensing Agreements
293. Use of Hard-Disk Drives for Storage of Sensitive Information
322. **Using on Company X Systems and Networks Prohibited**
329. Specific Configuration for System Availability
337. Dispersion of Computer and Communications Systems
338. Avoidance of Communication Network Central Point of Failure
340. Compliance with Standards Required for Emergency/Disaster Support
343. Five Category Application Criticality Classification Scheme
344. Preparation and Maintenance of Computer Emergency Response Plans
345. **Organization and Maintenance of Computer Emergency Response Team**
348. Preparation and Maintenance of Computer Disaster Recovery Plans
349. Preparation and Maintenance of Business Contingency Plans
352. Annual Inventory of Information Systems Hardware, Software, Etc.
354. Computer and Communications System Contingency Plan Testing
357. Access Control for End-User File Restoration Processes
360. Two Copies of Sensitive, Critical, or Valuable Information
363. Automatic Back-up to Local Area Network Server
414. Prior Approval Required for System Interconnection
415. Participation in Public Networks as Service Provider
417. Access Control Packages Required for Computers on Network
418. Formation of Binding Contracts via Electronic Systems
422. Large Networks Must Be Divided into Separate Domains
423. Dial-Up Connections Must Always Utilize Firewalls
424. Inter-Processor Commands from Outside Locations Prohibited
451. Modems on Workstations Connected to Internal Networks
452. In-Coming Dial-Up Lines Must Not Answer Until Fourth Ring
453. Maximum Permissible Password Attempts for Dial-Up Users
548. Designated Security Administrator for All Multi-User Systems

Figure 3.14 Local Area Network Security Policy List. (Reproduced with the permission of Charles Cresson Wood)



(c) 1993 Miura Systems Consulting, Inc.

All Rights Reserved

Business Resumption Planning

Figure 3.15 A flow chart of disaster recovery. (Reproduced with permission from Michael Miura.)

created by Michael Miora, one of the leading experts in this field. You will need to consider such things as "hot sites" that provide immediately accessible hardware and software that mirror your own when it has been rendered inaccessible due to disaster.

As with risk/business impact analysis and security policy drafting, DRP involves a lot of work. You might opt to contract out the entire task, from planning all the way through to hot site agreements and facilities. One company that can handle the whole process is Sungard Recovery Services. An alternative is to hire a consulting company to lead you through the process, such as Miora Consulting Services. Yet another option is to train your own staff to handle the job. Miora Consulting Services regularly conducts a very valuable one-day seminar entitled "Disaster Recovery Planning: The First Three Days." If you attend this seminar, you will be in a very good position to begin the DRP process within your organization.

Stay Tuned


One of the dangers of expensive and elaborate security systems is that, once installed, they give a false sense of security. In some respects, you cannot allow yourself to get too comfortable with the status quo. Remember that complacency is the curse of comfort. On the other hand, those responsible for security within an organization must avoid creating a constant state of paranoia, because this is as self-defeating as complacency. As Marshall McLuhan said, "The price of eternal vigilance is indifference." Issuing occasional security reminders to users and running the odd security awareness program should serve to keep attention to this problem at the right level.

The task of maintaining personal computer security is a cyclical one. The cycle of analysis and planning outlined previously really is only the beginning. Once equipment is in place, administration must see that it is used. Once rules are laid down, administration must see that they are adhered to. In larger organizations, it is wise to assign specific responsibility for these enforcement tasks, otherwise compliance will slide, a breach of security will occur, and there will be a lot of finger pointing without any useful decision as to who was to blame for the lax state of affairs. By giving final responsibility to one person, you probably will assure more diligent enforcement of security measures than if a group or committee is held accountable.

Once the ongoing enforcement of security policies is well in hand, it is time to consider beginning the cycle of planning and risk evaluation all over again. Certainly this is the case in large organizations. Circumstances change and so does equipment. The range of threats might increase. You have to assume that the sophistication of would-be intruders is steadily increasing. Revised assessment of risks, a review of current practices, and redefining of policy are all part of the ongoing task of ensuring personal computer security.

The Security Audit

So it is that we return to the information security audit, which is an in-depth review of current security status. This also can be the event that kicks off the entire security planning and assessment process. It also can be the undertaking that precedes that. The questionnaire in Figure 3.16 will give you an idea of the questions that need to be answered to get a clear picture of an organization's personal computer security

	<h2 style="margin: 0;">Workstation Review</h2>	Client Name: _____ Review Date: _____ Location: _____						
A. Equipment Details Workstation Name or Number: _____								
This workstation consists of the following equipment:								
Personal computer <input type="checkbox"/>	3.5 inch floppy disk drive(s) <input type="checkbox"/>	Mouse <input type="checkbox"/>						
Separate monitor <input type="checkbox"/>	5.25 inch floppy disk drives <input type="checkbox"/>	Other input device _____						
Printer <input type="checkbox"/>	Internal hard disk drive(s) <input type="checkbox"/>	Other output device _____						
Modem <input type="checkbox"/>	External disk drive(s) <input type="checkbox"/>							
B. Equipment Identification								
	Make	Model	Serial Number	Company I.D. Number				
System unit								
Monitor								
Printer								
Modem								
Other								
Other								
Other								
C. User Details								
Name of principal user: _____		Uses system for: _____						
Other users: _____		Uses system for: _____						
_____		Uses system for: _____						
Normal user hours: _____								
D. User Classification								
	1	2	3	4	5		Yes	No
Skill level of normal user	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Is there a keyboard lock?	<input type="checkbox"/>	<input type="checkbox"/>
Skill level of casual users	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Is the lock used regularly?	<input type="checkbox"/>	<input type="checkbox"/>
Security awareness of users	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Is boot security installed?	<input type="checkbox"/>	<input type="checkbox"/>
Regularity of backup	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Is any encryption used?	<input type="checkbox"/>	<input type="checkbox"/>
Sensitivity of data handled	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Is a modem used?	<input type="checkbox"/>	<input type="checkbox"/>
System exposure to public	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Locked drawer for backups?	<input type="checkbox"/>	<input type="checkbox"/>
Level of user consistency	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Is personal software in use?	<input type="checkbox"/>	<input type="checkbox"/>
E. Other Comments _____ _____ _____								

CQR Data Company, Form: SR0001 Revised 1/1/90 Copyright, 1990

Figure 3.16 Recording equipment details with FileMaker Pro

exposure. A good place to start is an inventory of personal computer resources. Most larger organizations have, or should have, centralized records of all personal computer equipment with serial numbers and current assignments.

The joys of being organized

In some cases, personal computers have been incorporated into the fabric of an organization so quickly that proper controls have been overlooked. Indeed, it is possible to find companies where staplers are more carefully accounted for than computers, software, and peripherals, all of which cost far more and open far greater avenues of potential loss. Considering the ease with which a simple database management program or electronic spreadsheet can handle the tracking of such details, there really is no excuse for not having a thorough record of all personal computer equipment. In Figure 3.17, you can see a simple database manager, FileMaker Pro from Claris, set up to record equipment details.

Among the hardware details to track for security, as well as general management purposes, are:

- Serial number
- Date of purchase
- Vendor
- Warranty period
- Maintenance contract
- Current location of equipment
- Current location of manual

Field	Value
Make	CompuShed
Model	286sx
Serial Number	C123245Y56
Purchase Date	12/9/90
Vendor	CompuShed, Inc (La Verne branch)
Warranty	Two years from purchase
Current Location	Sales office
Manual Location	Derek's drawer
Current User	Derek Smith

Figure 3.17 Data classification matrix

Keeping track of software is definitely a more challenging task, given the ease with which it can be installed/discarded by different users, the frequency with which upgrades appear, and the lack of consistency between vendors when it comes to upgrade methods. Nevertheless, a well-maintained database that reflects current software status will help immensely should there be a theft or security problems arising from software deficiencies. Among the software details to track for security, as well as general management purposes, are:

- Serial number
 - Original version number
 - Date of purchase
 - Vendor
 - Whether the software is registered
 - Current location of manual
 - Current version number
 - Date of last update

When preparing a security survey, you should consider the following questions on a station-by-station basis:

Is there a keyboard lock? This should be used.

Who has the keys? These should be controlled.

How skilled is the user? A low level of skill means a potential source of data loss; a high skill level could weaken effectiveness of simple defense.

- Does the software being used have undo capability? This should be turned on.
- Does the system have format capability? Remove it if practical or install protection.
 - Does the system have unformat capability? If not, install it.
 - Is a disk optimizing program installed/used? If not, install it.
 - Is personal software used? If so, is it guaranteed/approved?
 - How complex is the operating environment? The more complex, the greater the potential problems.
 - What type of backup device is installed? This is the first line of defense. Only diskless personal computers lack a backup device.
 - Is there an automated backup procedure? If not, install one.
- Is the backup media stored securely? If not, why do backup?
 - Is the backup performed regularly? Use backup file dates to check this.

One specialized but very serious window of vulnerability is any connection between a personal computer and other computers. This could be a local area network, a micro-mainframe link, or a modem for uploading and/or downloading data. The se-

curity exposure presented by such connections will be addressed in detail in chapters 11 through 13, but the important questions to ask are listed here:

- Is a modem attached? This is a potential window into the system.
- Does the modem accept incoming calls? This is an actual window into the system.
- Is the modem used for downloading data? It could be intercepted.
- Is downloaded data packed? Compressed data is harder to read.
- Is the data encrypted? This can ensure that intercepted data is not revealed.
- Is the modem used for downloading software? This is a possible carrier for viruses and Trojan horses.
- How reliable is the software source? Only use proven sources.

How personal computers are used

The most basic question to ask is, "What are the personal computers used for?" The answer to this allows you to determine what data is involved, what tasks will be interrupted by a security incident, and what aspects of the organization will be affected by data loss, misappropriation, or corruption. In some organizations, each personal computer has a different mission, so a system-by-system survey is necessary to answer this question. You might want to group systems by category of activity.

Receive information. Personal computers used for data entry, transaction recording, data acquisition, and even word processing are all fulfilling the role of receiving information. The security problems associated with such systems concern corruption or interception of incoming data and the disruption of processing capability.

Store information. The same computers might be the storage place for information, or they might hand it on to other systems. For example, a personal computer used for point-of-sale transactions might be the front-end of a retail accounting system and not retain much information for long. A word processor might receive a lot of information but store little of it beyond approval of the final document. However, it is important to bear in mind that many systems store quite a lot of data temporarily, even if they are not the data's final resting place. The security implications for systems that are used for storage center on the possible loss/theft/corruption of data.

Produce information. The primary task of some computers is to produce new information out of stored data. Such information might be new contracts assembled from boiler-plate text or revenue projections based on past performance and assumptions about the future. The security concerns for such systems are corruption of data due to errors and service interruptions, plus possible misappropriation of data for gain.

The impact of interruption

Whatever role a personal computer fulfills, continued fulfillment of that role is presumably of some value to the organization. You might have an old floppy disk PC that

is used for occasional word processing, the loss of which might not seriously affect business as usual, but the smooth and continuing operation of most personal computers is considered essential to mission fulfillment in most enterprises.

To examine the impact of an interruption in use due to virus infection, hardware loss or damage, or some other breach of security is to engage in "what-if." You have to use your imagination to list what ramifications an interruption might have. You can proceed on a station-by-station basis, assuming that each system in turn is rendered unavailable for one day.

The cost of loss

As an example of costing an incident, consider a supermarket chain that operates its own fleet of trucks and uses personal computers to manage the trucking operation. All maintenance information for the trucks is stored on a personal computer, including miles and hours logged, service performed, and so on. Vehicle maintenance schedules are produced and reports to management are generated detailing age and history of vehicles. Early one morning, the personal computer and printer are stolen. The backup disks from the day before are safe. Determine the cost of this incident:

Three people spend four hours each determining what has been stolen, arranging delivery of a rental unit, reporting the crime to the police, and briefing management on the problem: 12 hours.

Maintenance staff do not get their computer-generated work orders and have to figure out a schedule by hand, taking four people three hours: 12 hours.

One person spends three hours setting up the rental unit with software, backup data, and printer interface (rented printer not the same model as the one stolen): 3 hours.

Normal day is now over, but schedules need updating, daily entries made. Takes two people three hours at time and a half: 9 hours.

Insurance claim for loss of computer is prepared and pursued until paid: 8 hours.

Replacement computer is installed: 3 hours.

Cost of replacement computer and printer is \$6500 of which the insurance covers \$5000, leaving the company to cover a balance of \$1500.

Rental computer and printer returned: one week at \$200 per week.

The total cost of theft therefore is \$2170 (47 hours at \$10 per hour, plus \$200 for rental and \$1500 insurance shortfall). This is a fairly simple example, but it gives you an idea of how to proceed when evaluating the impact of a loss.

To make this example more interesting, you could add in the Fred Factor. Fred is a route driver delivering the supermarket's baked goods, which are sold to restaurants, as well as the bakery departments of the chain's own supermarkets. Fred earns a bonus for attracting new customers. He has spent weeks wooing a new restaurant that is near the end of his route. The restaurant owner's main concern is getting goods early enough in the day to meet the lunchtime rush. Fred has promised a first delivery the day of the computer theft.

Because of the disruption to the maintenance schedule caused by the theft, Fred's delivery van is not ready on time. Fred is late getting to the restaurant; the deal is off. Fred loses his bonus, the supermarket chain loses a sale that day, as well as a new chunk of business. While losses like this are hard to predict and to cost, they are very real. You can clearly see the importance of determining how critical each personal computer is to the organization's mission.

Classification of information

There are several ways to look at the importance of information to an organization. If you consider the impact of different forms of attack, this will help you realize the relative value of different types of information.

Damaged or destroyed. When data is lost and backup copies are not available, you can either recreate the data or live without it. In practice, you will find that a lot of disk space is taken up by files that are useful to have around but not important to on-going operations. A typical example is copies of correspondence in word processing files. Such files often are kept for reference purposes or in case similar letters need to be sent in the future. However, these letters probably exist in hard copy as well. If the files were deleted, it might be inconvenient, but the impact on the organization might well be minimal.

Accounting files present quite a different picture because recreating them can be very time-consuming. Furthermore, many organizations rely on such files in day-to-day decision-making. Without current data, operations could be severely impacted.

Altered or amended. The importance of data that is modified illicitly depends upon the extent to which the organization relies upon the data for operations and decision-making. If there is any possibility that it could be tampered with, data from personal computers that forms the basis of decision making must be verified before decisions are made. There must be safeguards against the tendency to assume that "It came off the computer, so it must be right."

Compromised or communicated. The importance of data that is disclosed against the wishes of an individual or organization depends upon how much damage the disclosure could cause. While losing your accounting files completely could bring your operations to an abrupt halt, having someone else reading them might have less immediate consequences. On the other hand, the word processing files that are relatively low in importance when destroyed might be high in importance when disclosed.

Indeed, this difference in emphasis can lead to damaging misjudgment. If a file that is of little importance to current operations is given minimal protection, this can lead to its disclosure, possibly with negative results. Like distilled liquors, files retained merely "for the record," actually can gain potency over time. Circumstances and opinions change, and material that was innocent enough when it was archived might prove to be incriminating a few years later.

Further classification

In order to decide how much security to accord the data that your systems handle, you can categorize or grade it. This will help set priorities if you cannot afford total protection for everything or if you find that according all data maximum protection would have a negative impact on productivity. You might want to use the following categories, bearing in mind that some data will fall into more than one category:

Decision-critical: Data that is used as the basis of decision-making (for example, budget projections, or current inventory).

Operational: Data that is an integral part of operations (for example, daily transaction accounts, time sheets).

Archival: Data that is retained for record-keeping purposes, such as purchase orders, invoices, and correspondence.

Convenience: Data that is retained because it might be useful at some point in the future, such as keeping expired contracts because the language might be used in new contracts.

Two more reasons for categorizing information are to raise user awareness of the importance of information and to protect the organization if something goes wrong. For example, if you label a floppy disk "Highly Confidential Data—Do Not Copy or Remove from Company Premises," there is not only a fair chance that the user of that disk will heed the warning but also a better than fair chance that you will be able to successfully claim in court that the user who sold the disk to your competitor should have known better. In Figure 3.18, you can see a data classification matrix, which is a useful device for grading information. This can be customized to meet the needs and priorities of your organization.

	Data Classification Matrix: Four Classes of Data, with Optional Categories			
	Sensitive	Confidential	Internal Use Only	Public
Loss, misuse, or unauthorized disclosure of this data could:	Have a serious negative impact and be very harmful to the company.	Have some negative impact and be harmful to the company.	Have little negative impact and cause only minor harm to the company.	Have no negative impact and would not be harmful to the company.
Accounting	<div>Categories</div>	<div>Classes</div>		
Marketing				
Personnel				
Manufacturing				
Administration				
Any data can be classified and optionally categorized by checking the appropriate box in this matrix.				

Figure 3.18 Recording survey data with a spreadsheet.

Application software used

Several questions of security can affect your choice of software. How reliable is the source of software? Is the software guaranteed to be virus-free? How strong are the security features built into the software? Does the software work with third-party security systems?

These questions are the first line of inquiry about the software upon which the organization is relying. A second round of questions should be considered. To what extent is the software customized? If so, who did the customizing? Are the persons reliable? Are they still around? Are they/their works bonded/insured?

Questions about customized software are particularly poignant given the prevalence of the "errant knight" syndrome in personal computer consulting. You will recognize this syndrome if you have ever encountered a personal computer that is supposed to perform a certain task and was "set up by this guy (the errant knight) who was a real whiz at these things." The personal computer is no longer performing the task correctly and this "whiz-kid" has long since whizzed off. These knights are errant in that they seem bound to wander and bound to leave a trail of used and confused users in their wake.

In fact, unscrupulous or merely unreliable consultants pose a very real security threat. Unless someone reliable within an organization has a handle on what is being done to and with the organization's personal computer resources, the organization can hardly rely on those systems performing correctly. For more on security exposure from software, see chapter 10.

The skill level of the users

This question requires more tact than most and is quite likely to bring a security survey into contact with personnel policies and organizational politics. Many users feel that they do not get enough training, while some managers feel that users "know too much already." There is something of a dilemma between a low skill level, at which accidents are more likely to happen, and high skill level, at which the ability to break simple security barriers is obtained. What you and your organization make of this dilemma will have a significant impact on security policy. Fortunately, some of the most straightforward security measures, such as file encryption and hardware locks, take exceptional skill to defeat.

Opinions vary greatly on the question of how much knowledge is a good thing. I happen to think that there is little to be gained by keeping people in the dark. For a start, the fewer secrets an organization has, the less it need worry about security at all. Secondly, time and the lessons of history oppose those who base their security on the ignorance of others. The benefits of personal computers are best realized by those who best understand them, and the more such persons you have in your organization, the more likely the organization is to succeed in its mission.

The organization that employs its own programmers might have a particularly tough time with the question of skill level and the following questions concerning ethics. For more on the personnel aspect of security, see chapter 14.

The ethical level of the users

Not many people will want to make the call on this subject, but it must be considered. Casting a suspicious eye over colleagues that you normally trust is no fun, but neither is losing data. There really are two levels of concern: the organization's defense against an attack by an unethical employee and the individual user's defense against an attack from within the organization. To what extent you trust those working alongside you will determine how keen you are to implement some of the hardware and file access control methods discussed in this book. To what extent the organization inspires and rewards loyalty will determine how far it has to go to protect itself from internal threats.

The hardware

Technically speaking, personal computer hardware never becomes obsolete. Sure, things break and people don't bother to fix them, but if you buy a personal computer to perform task A, then as long as it performs task A, it is not obsolete. New hardware might perform the task faster or run software that performs task B, which is all of A plus a whole lot more. However, there still might be a need for task A, and the old hardware can be kept in service long after its level of performance has been surpassed.

There is a lot of truth to this. The problem with using older hardware is not that it doesn't do the job, it is the difficulty that you will have replacing it if it goes bad. If your operations depend upon older hardware or purpose-built hardware, then you need to make sure that service, repair, and replacement remain available. Otherwise, you might lose the use of hardware in which valuable data is stored, effectively creating a technological tomb for your data.

Insurance

The insuring of personal computer resources is a fairly recent phenomenon. Insurance can be obtained at several levels and not all policies cover all of the losses that can accrue from a breach of security. Some other forms of insurance, like household contents, actually might exclude personal computers. Certainly a review of personal computer security should include a clear picture of what losses are insured, together with the records needed should it be necessary to file a claim. See chapter 4 for more on insurance and personal computer resources.

Current practices

Having evaluated an organization's personal computer security risks, assigning values to what might be lost, the process moves on to reviewing current practices. The goal is to get a clear picture of how existing security systems are used. While a survey will give you responses to questions, only by observing what actually happens in day-to-day operations will you get an honest picture of current exposure.

For example, if it appears from observation or spot checks that employees conscientiously make daily backups that they store in a fireproof safe at the end of the day,

then you have less to worry about than if you find backup media lying around on desks after hours. Indeed, a walk around the office when everyone else has gone home can be very revealing. If you see a lot of floppy disks left on desks, keys left in locks, and systems left turned on and unprotected, then you know that current practices leave a lot to be desired.

The story is told of a consultant who was having great difficulty interesting a potentially lucrative client in his security services. After being rebuffed several times, he scheduled a last-ditch presentation late one Friday afternoon. On the pretext of visiting the bathroom, he managed to make a quick unaccompanied tour of the premises. When he returned to the conference room, he began his presentation by showing a handful of keys. His audience finally was convinced that there was a security problem when the consultant revealed that these were the keys to the company's personal computers, left in the locks by employees who had left for the weekend.

Methods in general

For a larger company, the tasks of assessing security is more complex but usually can be accomplished with an expanded version of the approach outlined in the previous example. Probably the best place to start in a large organization is with a station-by-station review of the current deployment of personal computers. Until you have a clear picture of what each personal computer does, you cannot formulate effective strategies. You can use a personal computer to assist you in this survey. Earlier, in Figure 3.16, you saw a sample questionnaire laser-printed from a PC.

Once the questionnaires are completed, the responses can be compiled and analyzed with software. Most database or spreadsheet programs can be adapted to the task of collecting answers. Some basic statistics then can be developed, such as the total number of stations, the percentage that have modems, and so on. Reports that list comments can be printed out for review. As a whole, the data collected will give you a detailed picture of the current situation. You then can draw up a list of potential threats and begin to assess the degree of risk posed by each one. With survey responses about the type of data processed and basic figures on the organization's revenues and expenditures, you can begin to assign values to data, equipment, and potential losses. In Figure 3.19, you can see an Excel spreadsheet designed to accept survey responses keyed in by a data entry operator.

Software assistance

What type of software you use to collate your information will depend upon what you are used to and what reporting capabilities it has. For example, the FileMaker Pro database manager works on both Macs and Windows machines and allows you to set up very friendly data-entry forms like the one shown earlier in Figure 3.17. Indeed, you might want to set up a live questionnaire on computer. A disk with the questionnaire program could be distributed to employees, allowing them to enter their responses "online." If your organization's personal computers are networked, the survey could be carried out over the network. Alternatively, you could make the rounds interviewing employees face-to-face but entering their responses directly into a database on a portable computer.

Microsoft Excel - SECREV1.XLS						
File Edit View Insert Format Tools Data Window Help						
	A	B	C	D	E	F
1	Workstation Review					
2	Client Name: ConCerned Industrial			Please complete all sections as		
3	Review Date: 11/15/95			accurately as possible		
4	Location: San Vinto			Thank You!		
5	A. Equipment Details					
6	Workstation Name or Number: Sales #2					
7	Components	Y/N	Type	No.	Components	Describe
8	System unit	y	3 5 inch floppy	1	Mouse	Msoft
9	Separate monitor	y	5 25 inch floppy	1	Other input	No
10	Printer	y	Internal hard drive	1	Other output	No
11	Modem	n	External hard drive	0		
12	B. Equipment Identification					
13		Make	Model	Serial Number	Company ID	
14	System unit	Dell	486SX	D2345235L2	S2001	
15	Monitor	Dell	VGA Plus	DJK283838	S2002	
16	Printer	HF	LaserJet	452A343	S2003	
17	Modem		No			

Figure 3.19 Questionnaire form on computer

As with other aspects of the entire audit-analyze-plan cycle, there is software assistance available. A number of network auditing tools are discussed in chapter 11. These allow you to ascertain the hardware and software status of network workstations as well as test for security standard. Blue Ocean software makes a program called Track-It, which is designed to automate the tasks of recording hardware and software details. In Figures 3.20 through 3.22, you can see this program at work. In Figure 3.20, you can see the basic information about a user and the software on her machine. Pressing the F6 key brings up details of her hardware configuration, as shown in Figure 3.21, whereas F7 displays network connection information, shown in Figure 3.22.

The Network Connection

Organizations that use local or wide area networks must establish a security policy and plan for disaster recovery or face the prospect of going out of business. This might sound drastic, but it is a fact that networks are much harder to resurrect when things go wrong than mainframe systems. Average down time for LANs is longer than for mainframe systems, and the longer a network is down, the greater the risk that the impact on the organization will be fatal rather than merely damaging.

If you use your LAN for anything more serious than playing DOOM and you don't have policies and recovery plans in place, I urge you to make these priority action

(c) 1993 BOSI		T R A C K - I T !		v2.0	
NAME :	BRYAN, LAURA	WS # :	5		
DEPT :	FINANCE	DEPT # :	200		
PHONE :	(555) 555-5555	EXT. :	1234		
COMPAQ 386-200		S/N: 94671240			
HARVARD GRAPHICS		S/N: 985478478			
LOTUS 1-2-3		S/N: 3497237			
MICROSOFT MOUSE		S/N: 672372398			
MICROSOFT WINDOWS		S/N: 849578			
NEC MULTISYNC 3D		S/N: 2378234978			
F5=USER INFO F6=CONFIG F7=LAN F8=MARK ITEM F9=TRANSFER ITEM F1=HELP					

Figure 3.20 Using Track-It to record details of hardware and software

(c) 1993 BOSI		T R A C K - I T !		v2.0	
NAME :	BRYAN, LAURA	WS # :	5		
DEPT :	FINANCE	DEPT # :	200		
PHONE :	(555) 555-5555	EXT. :	1234		
TRACK-IT! AUDIT RESULTS (01/07/92 at: 14:50:59)					
CPU :	80386SX	DOS VERSION :	5.00		
CPU SPEED Mhz :	20	BUS :	ISA		
MATH CO :	NONE	ROM BIOS DATE :	04/09/90		
MOUSE TYPE :	INPORT	PARALLEL PORTS :	1		
MOUSE VERSION :	2.4	SERIAL PORTS :	1		
KEYBOARD :	101	GAME PORTS :	0		
VIDEO :	VGA				
VIDEO RAM :	256	DRIVE	SIZE(MB)	FREE(MB)	
DOS MEMORY :	640	A:	(1.2 MB 80 track 5.25)		
FREE MEMORY :	571	B:	(1.44 MB 80 track 3.50)		
EXT. MEMORY :	1024	C:	68.10	0.96	
EXP. MEMORY :	0				
ESC = QUIT SCREEN C O N F I G U R A T I O N F10 = SAVE SCREEN					

Figure 3.21 Configuration information recorded by Track-It

items. If any of your systems have Internet connections, then such action is even more urgent. The best place to begin is with a security audit, which can be performed inhouse or by hiring specialists. While NCSA can arrange such audits, you might find that your accounting firm also offers such services, or you might have a word-of-mouth recommendation from fellow professionals in organizations such as IEEE (Institute of Electrical and Electronics Engineers), ACM (Association for Computing Machinery), or DPMA (Data Processing Management Association). If you are interviewing prospective consultants, ask if they are CISSP qualified (as in

(c) 1993 BOSI		T R A C K - I T !		v2.0	
NAME :	BRYAN, LAURA	WS # :	5		
DEPT :	FINANCE	DEPT # :	200		
PHONE :	(555) 555-5555	EXT. :	1234		
TRACK-IT! AUDIT RESULTS (01/07/92 at: 14:50:59)					
LOGIN NAME :	SUPER				
NET ADDRESS :	05F06				
NODE ADDRESS :	0000790014F5				
NET/XMS/EMS :	XMSNET				
IPX VERSION :	3.10				
SPX VERSION :	3.10				
DRIVER VERSION :	1.10				
SHELL VERSION :	3.22A				
LAN CARD :	Networth Inc. - EtherNext/16 V1.10EC (901202)				
INTERRUPT :	15				
I/O BASE(HEX) :	340h				
ESC = QUIT SCREEN		L A N I N F O		F10 = SAVE SCREEN	

Figure 3.22 Network details recorded by Track-It

Certified Information Systems Security Professional, which is a professional accreditation program administered by ISO²—International Information Systems Security Certification Consortium). You can contact all of these organizations via the W.E.B. page listed at the beginning of this chapter.

Summary

Policies, plans, and audits take time to create. This chapter has suggested numerous resources to help you speed the process, such as specialized software and outside consultants. However, just about any organization can find within itself enough of the vital ingredients of security planning: imagination, clear thinking, and common sense. Apply these to the following questions, and your organization will be better equipped to gain, rather than lose, from its investment in personal computers:

- What could go wrong?
- What would the impact be?
- What should be done to minimize B and prevent A?
- How should we react if A happens anyway?

HAVING JUST COMPLETED THE ANNUAL BUDGET, HE THOUGHT
HE HEARD THE WELCOME SOUND OF EDNA'S TEA TROLLEY.



"Used by kind permission of Paul Davies Publications."

Secure Hardware Defending and Insuring Equipment

W.E.B. connection: <http://www.ncsa.com/pclan/chap04.html>

*"Tis in my memory locked,
And you yourself shall keep the key of it "*
SHAKESPEARE, HAMLET, ACT: I, SCENE: III

*"The computer is America's second most
frequently stolen item "*
LIGHTGARD PRODUCT BROCHURE

This chapter concentrates on preventing hardware from being stolen, vandalized, or otherwise damaged. This does not cover the question of maintaining a smooth supply of power, which is covered in chapter 5, or how to stop unauthorized people from entering your office and using your hardware, which is covered in chapter 6. If you want to know how to stop people from using electronic eavesdropping techniques to illicitly obtaining information from your hardware, see chapter 7.

Although removable media such as floppy disks are not considered hardware, they do have a tendency to go missing. So this chapter addresses the problems of stopping such items from "going walkabouts." Because there will always be limits to the effectiveness of whatever protective measures that you take against hazards such as theft, power failure, and lightning, this chapter also provides tips on insuring your hardware.

Physical Security

Protecting your hardware from disappearing is known as *physical* security. In this section, I will discuss, in general terms, what physical security involves, then map out an approach to the subject in the rest of this chapter.

The scale of the problem

Just in case you need reminding that physical security is a serious problem, take a look at the charts in Figures 4.1 and 4.2. The pie chart in Figure 4.1 shows that theft accounted for half of all equipment losses in 1993, exceeding the combined losses due to fire, water, surges, and accidents. These figures are nationalized estimates calculated by SAFEWARE, the Ohio-based computer insurance specialist, based on reported claims. To give you an idea of the scale of losses and the dramatic rate at which they have increased, Figure 4.2 compares losses in 1986 with those in 1993. It is interesting to note that SAFEWARE pegged total 1993 losses, from all causes, at \$1.3 billion. In 1993 losses due to theft alone exceeded \$1 billion.

Obviously, there are many possible explanations for such a rapid rise in losses from theft. Three factors are particularly important as they require a shift in security assumptions compared to just a few years ago:

Computers have street value. Most people in Europe and America use, or have access to, a computer. Most people can put a price on personal computer equipment, and most people want a computer. This places computers in the same category as stereos, VCRs, and TVs. (In 1994, Americans spent as much on personal computers as they did on television sets.)

Criminals are computer literate. They know the relative value of different computer components. There are plenty of examples of "stealing to order" in which thieves work to a shopping list of specialized hardware.

Organized crime is involved. Many computer thefts are petty affairs involving users of illegal drugs desperate to finance their next purchase, but an increasing number involve organized crime directly. This includes supplying computer equipment to dollar-starved organizations in Eastern European countries or acquiring the means to convert criminal profits into nontraceable goods with a universal cash value (such as RAM chips).




Figure 4.1 Pie chart showing value of losses in 1993 according to category (Figures provided by SAFEWARE)




Figure 4.2 Column chart comparing losses in 1993 with those in 1986 (*Figures provided by SAFEWARE*)

The physical territory

For most organizations, the first line of defense against equipment theft is not physical but psychological. In other words, we rely on the tendency of most people to behave, at most times, in an honest and ethical manner. Yet even when you trust all of the people who normally would be in a position to steal your computer equipment, you still have to consider those other people, the ones who are prepared to break in and steal



Detectives Break Racket

Brisbane crime squad detectives have cracked a major breaking and entering racket that netted stolen computer equipment valued at more than \$150,000.

Australian Associated Press, March, 1995

The second line of defense thus is site or perimeter security. Because site security is something that every organization and individual is concerned about, even those who don't use computers, this book only addresses the more advanced aspects, such as smart cards and biometrics. Because these technologies are used to control computer access as well as building access, they are discussed in chapter 6.

The third line of defense is restraint, as in making it difficult for people to simply pick up equipment and walk off with it. Following on from this are alarm systems that let you know when equipment is being moved. Because some of today's thieves are computer-savvy, you also have to defend internal components such as memory chips, which often make a more attractive target than complete systems. A good general deterrent to theft is proper identification of property. I will look at some sys-

tems that provide this specifically for computers. In addition, I will consider several examples or case studies that demonstrate the need for, and the practice of, physical security.

The physical risks

Physical security is not just about avoiding the cost of replacing expensive equipment. If physical security is weak, you invite risks in all three areas of computer security:

Confidentiality

Integrity

Availability

Almost all of today's personal computers contain data, which typically is stored on a hard disk. Although many of the people who steal computers are interested in the street value of the hardware rather than the data contained therein, you have no way of knowing this. Consequently, such theft represents a major breach of confidentiality. If someone steals a computer from your office, you might well have to face the fact that sensitive data now is in the hands of people whose ethics are clearly in doubt. There are steps that you can take to defend confidentiality under adverse circumstances, such as encrypting data so that a thief cannot read it, but the focus of this chapter is preventing the theft from happening in the first place.



A Sick Case

Employees found three computers and several diskettes missing from the South Florida AIDS Network offices at Jackson three weeks ago. ~~he~~ hard drives and diskettes contained highly confidential information on services received by at least 6000 to 7000 people who get assistance through the publicly funded agency. Health officials stressed that the information on the computers' hard drives and diskettes was protected by several passwords and access codes. Investigators said Monday they were questioning staff at the hospital, including custodians and supervisors, who might have had access to keys to the locked offices where the computers were kept. There were no signs of forced entry.

The *Miami Herald*, December, 1993

Theft poses several possible threats to information integrity. If a computer is stolen and later retrieved, how much confidence can you have in the integrity of the data that it contains? Even if the thief did not mess with the data, physical damage is a real possibility.

Theft can seriously affect information availability, particularly if you don't have backups of the information stored on the stolen equipment or if the backups also are stolen. However, even if you have backups of your information, that might not be a complete defense against the threat to availability. If all of your computers are stolen, there is bound to be "downtime" until backups can be restored onto replace-

ment hardware that you have rented or purchased. If you have a lot of computers in your office, you might think it is unlikely that anyone could make off with all of them at once. You might be right, but thieves could render all of them unavailable. One person can carry 1000 computers' worth of RAM chips.

Sources and resources

As you saw in chapter 2, physical security is largely a matter of common sense; however, in the past, it was difficult to find all of the pieces that you required to create a complete physical security solution for computer equipment. Today there are plenty of products available. You will find many of them listed in the Physical Security section of the *Infosecurity News Buyers Guide*, which is on the Cobb/NCSA Security Resource Disk included with this book (see Figure 4.3). The section is divided into three parts:

- Equipment anti-theft and component-locking systems
- Facility and entrance controls
- Laptop security

Assessing physical security

Writing in the 1989/190 *Information Security Guide*, security consultant Gerry Faulks describes the four-fold role of physical security:

- To deter
- To delay

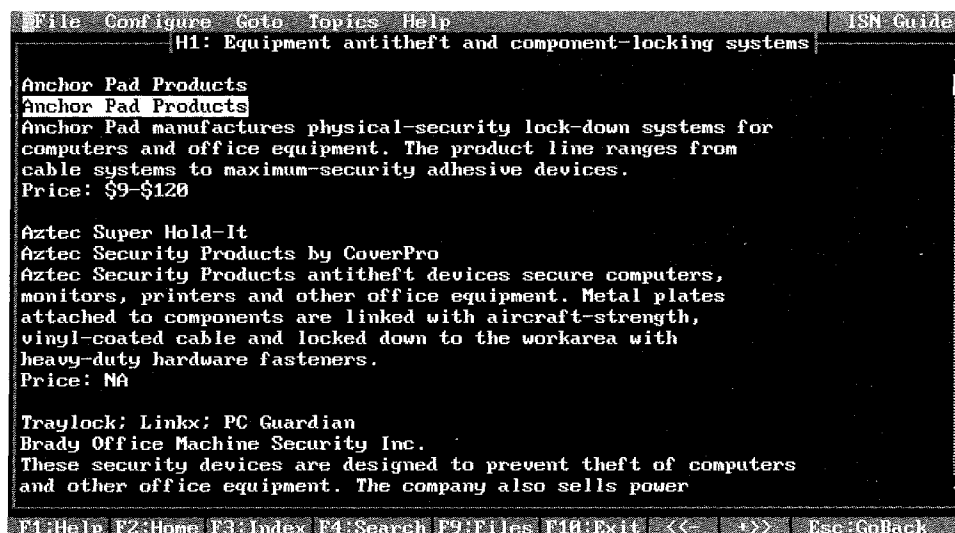


Figure 4.3 The first few Physical Security listings in *Infosecurity News Buyers Guide*

To detect

To defend

As different strategies are reviewed in this chapter, their performance in these four areas will be assessed. Making sure that your personal computer hardware stays where it should can be as simple as securing the hardware to the desk. Yet even simple solutions need to be thought out. They can be effectively implemented only after careful analysis of the potential threat. For example, attaching an expensive personal computer and printer to a computer work desk might be of limited value if the work desk has wheels and is located near the freight elevator.



Facts of Life

According to Laurence Milledge of Coventry-based security specialists LRM Consultants, "in today's office environment, theft prevention is all about response time." In other words, you have to base your security plans on the fact that, if something has a street value, people are going to try and steal it. Milledge notes that "Alarms alone are not sufficient deterrent for today's thieves, who assume that, when they break in, alarms will go off." Consequently, they "concentrate on stuff that is easy to pick up and carry off."

A Secure Example

For several years, beginning in 1985, I taught seminars at IBM offices in a large high-rise building in downtown San Francisco. The security arrangements in use by IBM over a decade ago give you some indication of how seriously hardware security is regarded by the world's largest computer maker.

The classroom was on the third floor, and all elevators that accessed that floor were monitored by security personnel. Only badge holders were allowed onto the elevators. This meant that all seminar attendees had to sign in with a security officer and get a badge before going up to the classroom. You had to be preregistered for the class, and the security officer had a list of preregistered attendees. When you reached the third floor, there were two doors. One led to the area that contained the classroom, and the other led to IBM's internal computer facilities. This second door was controlled by a magnetic card reader and was operated by employees with the correct ID cards (that's ID as in identity).

In the classroom itself, the PCs were set out on proper workstation desks. All of the monitors and system units were attached to the desks by cables anchored on the underside of the work surface. The cables were connected to the monitors and system units with adhesive pads. The monitors could be moved for better viewing and the system units could be opened for quick repairs, but neither could be removed from the room without the use of a screwdriver (see Figure 4.4).

When students took a break, they were issued with special cards that enabled them to unlock the door leading to the section of the floor containing the bathrooms. These cards did not unlock any other doors. All cards were collected at the end of the class. This arrangement was not entirely convenient, but it was secure. Only a




Figure 4.4 A cable restraint system (in this case the Kablit KAB 100 from Secure-It). (Photograph used by kind permission of Secure-It)

determined intruder would persist in face of the security barriers that IBM had erected. Whether you think you need to go to such lengths or not, the best frame of mind in which to read this chapter is a suspicious one, reflecting on the ways in which your personal computer hardware could be ripped off, and the best defense against that eventuality.



Taking Computers from Kids

In a brazen daylight heist Saturday and Sunday, police said three to four suspects scaled the roof of the three-story Westville Street School, kicking in a window to gain entry. Eleven computers, two laser printers, a satellite receiver, a fax machine, a video cassette recorder, a CD disc player, and two telephones were stolen, school officials said, adding that none of the merchandise was insured. Thieves also stole all of the computer discs, which contained a year's worth of work from 120 fourth- and fifth-grade students.

Boston Globe, April, 1993

Exercising Restraint

One of the most straightforward approaches to preventing desktop equipment from disappearing is to attach it to the desk upon which it is used. Preferably this will be

a fairly large piece of furniture, heavy, awkward, or otherwise difficult to move. The idea is that equipment is much harder to steal if it requires that the furniture goes with it or if it first has to be detached from the furniture. There are roughly three categories of restraining device.

Basic security

The most basic restraint systems literally attach to the base of the secured equipment. For example, the Securit system from LRM Consultants consists of interlocking steel bars, the top halves of which are attached in pairs to the underside of equipment such as printers, monitors, and system units. This is done with existing bolts or industrial adhesive. You attach the bottom half of the bars to the desk, again with bolts or adhesive. The equipment then "docks" into place and is secured with a hard-to-pick rotary pin tumbler. The bars form solid interlocked units offering no point of access for crowbars. If bolts are used for the attachment, the nuts are hidden within the bars when they are locked in place (see Figure 4.5).

Removal of the equipment for service access is simple: Turn the key and lift it off. The bar is designed to give solid level footing even when the unit is not docked. The only administration involved in the system is looking after the keys. This system is highly effective for a wide range of equipment, including VCRs, TVs, and laboratory




Figure 4.5 The Securit locking bar system from LRM Consultants. (Photograph used by permission of LRM Consultants)

devices such as oscilloscopes. However, you already can see where some of the trade-offs occur with hardware restraint systems. Some people might have doubts about the aesthetics, but I will assume that you would prefer to look at your equipment secured with a restraint system rather than an empty space where it used to be.

The main practical disadvantage of Securit and other devices of this kind—such as the Lucaskey Tray-Lock, the Compu-Gard range, and the AnchorPad device shown in Figure 4.6—is that they impose restrictions on tasks such as rearranging the office. Are you going to end up with holes in the desk that are no longer needed? Also, you need to make sure that users can adjust the positioning of their monitors; otherwise, your security devices might have a negative effect on ergonomics, which are in turn a "risk" factor for organizations that rely on computers. The same applies to systems such as the all-in-one Apple Macintosh and Compaq Presario models, where the monitor is integral to the system unit.

In common with the Anchor Pad system in Figure 4.6, some devices include a casing around the system unit to which the monitor can be attached. However, caution should be exercised when using locking or docking devices to secure monitors. The swivel bases on some monitors are not securely attached to the main body of the monitor and thus do not make a good anchor point. In summary, here's how this type of protection fares in four categories:

Deters: Probably. A good deterrent to most thieves. Not a deterrent to vandals.

Delays: Definitely. Getting away with the equipment is a much lengthier procedure, even for the experienced thief.

Detects: No. While the restraining device might show signs of tampering, it has no inherent ability to tell you who did the tampering and so cannot contribute much to the detection of thieves, would-be or successful.

Responds: No. These are simple mechanical devices, which makes them very cost effective, but they do nothing to sound the alarm.



Read All About It

More and more physical security systems are appearing on the market to meet the growing concern over equipment theft. To stay abreast of these developments, be sure to subscribe to industry publications such as *Infosecurity News*. Because many physical security products have applications beyond the specific field of computer security, you also should subscribe to a general security trade journal such as *Security*.

Getting cable

An alternative system of restraint is based on cables and locks, reminiscent of bicycle locks and some antishoplifting devices. This system is available in several forms from several companies, such as the Technalock system from Business Security Systems and the Kablit system from Secure-It. Clearly there are aesthetic penalties involved when you add restraining cables to the already annoying clutter of cords and




Figure 4.6 The Anchor Pad system in action, securing a Macintosh PowerPC

cables that most personal computer systems require. However, the system provides more freedom of movement than the anchor devices. Most cable-to-desk systems are made up of one or more of the following components.

Cable. A steel cable is passed through an existing opening in the hardware or through an eyelet attached to the hardware with a sticky plate or bolts. The cable is attached to the desk with screws, bolts, or a sticky plate, as shown in Figure 4.7. To prevent scratching from bare metal, some companies supply plastic-coated cable. This also reduces the possibility of electrical problems, such as short-circuiting.

Sticky plates. When hardware has no obvious place to which a lock can be attached or through which a cable can be passed, you can create one by using a large area of adhesive, as shown in Figure 4.7. The industrial adhesive used on these plates is pretty effective, withstanding thousands of pounds of force. However, you might not want your equipment permanently marred by such additions. You also can use sticky plates to terminate cable on the back or underside of a desk.

Concealed bolts. To attach cables or loops to a piece of hardware, you can use surface-mounted screws or bolts that already are holding things together. You remove the screw, pass it through a cable socket, then replace it. When the cable is passed through the socket, the screw is inaccessible, as shown in Figure 4.7. An ingenious variation, called a *hingefastener*, allows you to use the bolts on the underside of the system casing. Yet another option is the socket shroud, which is shown in Figure 4.7. However, you should always exercise caution about bolts and screws that you press into service for restraining devices. They might not hold things together as well after you give them the added task of holding a cable socket. Concealed bolts can be used to terminate cables on desks.

Locks and keys. The restraining cable needs to be secured at both ends, with at least one end capable of being disconnected to allow removal of items that are secured by the cable. There are a variety of ways to accomplish this. Many systems use a special form of padlock into which a cable end can be inserted and locked. The cable can be looped around the desk frame and locked or threaded through a hole in




Figure 4.7 Diagram detailing cable restraint system

the desk with a lock on the end to prevent removal. Some options are illustrated in Figure 4.8.

You don't need locks at both ends. One end can be arranged so that it cannot be removed unless the other end is unlocked. The locks that you use can be resettable combination locks or keyed padlocks. See the later section "Key management" for more about selecting and using a set of keys.

The Mac and other factors

For many years, Apple has been building "security retention sockets" into the casings of Macintosh computers and other devices, such as printers and external drives. These allow you to insert a tab that can be locked into place and through which restraining cables can be passed. This system is a positive innovation because it makes physical security capability part of the hardware itself, rather than a completely separate "add-on." Other manufacturers would do well to follow Apple's lead in this area.

Unfortunately, there are other aspects of physical security where Apple's example is not good. The 6100 series PowerPC Macintosh models are notorious for two design flaws. The first is the power on/off button, which is located right where a PC user would expect the floppy disk eject button to be, resulting in numerous accidental power downs, with attendant loss of unsaved work, until the user gets used to the fact that Macs have no disk eject button. The second is the lack of any retaining screws for the system unit lid.

The 6100 system unit has been dubbed "the candy box" because you can pop it off with absolutely no tools, exposing an array of valuable internal components for the thief to choose. Even if there had been just one retaining screw, as on the first "system unit" style Macintosh (the Mac II), users concerned about theft could use it to attach a cable socket and thus seriously impede efforts to open the case. There is a built-in security socket that helps secure the unit to the desk, but many of today's




Figure 4.8 Options for securing cable ends

thieves are smart enough to pick and choose the small, high-priced internal components rather than bother with bulky system units. (The Anchor Pad casing shown in Figure 4.6 should solve the problem.)

You might want to bear in mind several nonsecurity factors when using cable systems. Most cable/connect systems are lacking in aesthetic appeal. At first sight, they tend to look cumbersome and give off an air of paranoia. To some people, they are reminiscent of pictures in a motel room that are bolted to the wall. However, users quickly grow accustomed to them. Note that restraint systems alone do not prevent damage such as vandalism. The inconvenience factor also should be considered. Servicing and moving components is made more complicated. When padlocks are used, keys or combinations need to be managed. When evaluating a cable restraint system, you need to decide how it performs in the four categories.

Deters: Probably. A good deterrent to the casual thief who is not familiar with the restraining device. Less of a deterrent to the employee or other person with the time and means to unscrew/unstick/unlock.

Delays: Definitely. Makes running off with a PC a lengthier procedure, however adept the thief.

Detects: No. While the restraining device might show signs of tampering, it has no inherent ability to tell you who did the tampering and so cannot contribute much to the detection of thieves, whether would-be or successful.

Responds: Possibly. Some restraining devices can be connected to alarms, setting them off if tampered with. The response to the alarm depends upon the sophistication of the device and the priority assigned to the alarm.

A strong case

Special attention must be paid to equipment used in areas where there might be few limits on site access, such as government offices, schools, hospitals, shops, and building sites. In such cases, securing computers to the furniture might not be enough. You might feel safer locking the entire unit away when it is not in use. Several companies sell products designed for the safe storage of computers, word processors, and similar valuable equipment.

Features to look for include continuously welded, fully rebated steel plate construction, completely enclosed locking mechanisms, heavy-duty high-security locks, antileverage channels, secure steel doors or shutters, cable organizers and electrical outlets, heavy-duty castors, and adjustable shelves.

So, how does the practice of locking up a personal computer system in a steel case perform against the physical security criteria?

Deters: Yes. A very good deterrent to all but the most determined thief.

Delays: Definitely. If the thieves set off an alarm when they broke into the building, this type of protection might well cause them to pass over the protected item.

Detects: Possibly. Attempts to force the unit open might leave visible evidence.

Responds: Possibly. It would not be difficult to fit an alarm to the case that would let you know if it was being attacked.



Lesson Learned?

Thieves may have a hard time unloading their loot from a recent heist. They stole a computer system that teaches police officers when to shoot at suspects. Video Training, Inc. of St. Louis estimated that the computer system taken from a trailer was worth about \$50,000. According to company President Jack Kootman, the thieves got all the components and tapes, leaving behind only the screen.

St. Louis Post-Dispatch, March, 1993

Portables, Alarms, and Other Ideas

Portable computers present special problems as they are easier to carry off than desktop models, have a higher per-pound street value, and often travel outside the relatively secure confines of the office. Alarm systems can be fitted to both portable and desktop machines, and a variety of other ingenious physical security systems have been developed.

Power cables

A clever twist on antitheft cables is found in the LightGard product from Interactive Technologies and the Phazer system from Computer Security Products, both of which use fiber-optic technology. Instead of cable that relies on strength, this cable, which is very thin and easy to thread through sockets and brackets, carries a continuously monitored optical signal. The cable is arranged in a loop and creates a live fiber-optic circuit. The black box into which both ends of the loop are plugged transmits a pulsating beam of light. Any attempt to tamper with the cable sets off an alarm and pinpoints the location of the trouble (see Figure 4.9).

The black box can transmit the alarm with a radio signal to a separate monitoring station or can be hardwired to any existing 12-volt alarm control panel. Because the system does not rely on fixed attachment of the cable to the equipment, it offers somewhat greater freedom of movement than steel cable systems. Individual stations can be bypassed out of the system for servicing or relocation. How effective is the system? Carnegie-Mellon University reports that, during 1992, some \$250,000 worth of unsecured equipment was stolen, but there were no thefts of the 3000 items secured with LightGard.

Deters: Definitely. Extensive inside knowledge or sophisticated equipment is required to defeat the system.

Delays: No. Cable is easy to cut or unplug.

Detects: Yes. The system pinpoints where tampering is taking place.

Responds: Any tampering immediately sets off an alarm.

Alarm and power protection

Battery-powered alarm systems are available to fit inside or outside desktop units and notebooks. Some devices, such as the SonicPro from SonicPro International, act as motion detectors, sounding an alarm if equipment is moved or disturbed. Other




Figure 4.9 A fiber optic cable restraint system

systems, such as the Squealer from Stajer Corp., plug into an ac outlet and, if the current is cut or the system is removed, sound an alarm.

Business Security Systems sells a universal device to prevent the unauthorized "plugging-in" of any equipment with American-style plugs. This is a combination lock that covers the prongs of the plug. While this would not be much of a deterrent on equipment that uses removable power cords, which includes most system units, it is one more item to include in your armory.

The buddy system

A software approach to hardware security can be found in Security Force from Globus Systems and other "buddy" systems. These work in networked environments. Each network computer is assigned a nearby "buddy." The buddies watch over each other at all times (typically by sending and receiving status messages across the network). If a protected computer disappears from the network without proper authorization, its buddy transmits an alarm to deter the thief and help detect the theft in progress. The system also is an effective deterrent to theft of internal components—such as boards, chips, and disk drives—as it is almost impossible to remove these without changing the network status of the computer.

Drive locking

While desktop computers usually have screws or bolts that can be pressed into service for the attachment of restraining devices, many portable systems do not. One way of addressing this problem is to use a drive lock. This is a device that is inserted

into a floppy disk drive and then twisted and locked in place so that it cannot be removed. These devices are discussed in chapter 6 from the perspective of system access control, because they stop people from putting floppy disks in the drive. From the physical security perspective they offer a means of attaching computers to a restraining cable, as illustrated in Figure 4.10.

Another example of portable computer restraint can be seen in Figure 4.11. However, while such devices are an excellent idea when using a portable system in an office setting, they don't protect laptops when they are at their most vulnerable, which happens to be when they are most portable. Theft of portable computers from parked cars, hotel lobbies, and airport lounges is an occupational hazard for business travelers who need to take their computers with them.

Fortunately, there are some simple precautions that you can take. First of all, ask yourself if you really need to carry your computer in a portable-computer bag. They might look nice, but they practically shout "Come and get it." I like to tote my Compaq Concerto in an ordinary leather briefcase of the type that has a shoulder strap. Even at close quarters, you cannot be sure that there is a computer inside. Because many thefts of laptops occur in a "target rich" environment, the perpetrator is likely to pass on any targets that are ambiguous and go for the sure things (for example, those bags with "Toshiba" or "Compaq" emblazoned on the side).

Other precautions include obvious, generic anti-theft measures such as not leaving the computer in your car or hotel room (unless you use the in-room safe). When leaving your laptop at the office overnight, put it in a locked drawer. Make sure you




Figure 4.10 The Secure-It SEN 600 floppy drive lock for portable computers




Figure 4.11 Lapkit from Aztec Security Products.

have insurance coverage for the laptop, even when it is in transit (you cannot rely on car-contents, house-contents, and office-contents insurance for this—specialized coverage is available from SAFEWARE and others). Make sure you keep the machine backed up (parallel port tape drives, like the Iomega Ditto Easy 800, are handy for this) and consider an access-control program that encrypts the drive (see chapter 7). This renders the data on the drive inaccessible if it falls into the wrong hands.

Several alarm devices now are available for both portable and desktop computers. Some are triggered when the computer is picked up, while others go off when the computer is moved out of range of a tracking device. I have not yet had any feedback on the effectiveness of these alarms, but they might merit consideration. One obvi-



Online Help

One way of checking up on physical security solutions is to visit the NCSA Information Security Forum on CompuServe. There, you can ask other users, managers, and security professionals for their opinions on specific applications. For example, when one user asked about securing a new shipment of Macintosh PowerPC machines, he received half a dozen suggestions, including pricing and the names and addresses of suppliers.

ous limitation is the fact that, once someone has decided to walk off with your laptop, the fact that it is sounding an alarm might not slow them down (in fact, they might break into a run). While there is a good detection factor, there is only a limited deterrence and far less prevention than a restraint system.

Tamper Resistance and Identification

Now that personal computer components have "street value" just about anywhere in the world, a locked system case is becoming a minimum security requirement. It has the added benefit of discouraging would-be technical experts from interfering where they shouldn't. Some manufacturers now provide rotary key case locks as standard equipment, but others do not. Cable systems and security casings can be used to compensate for the lack of case locks, but you should demand integral case locks as standard equipment. Fortunately, brute force and high technology are not your only weapons in the effort to establish and maintain physical security; you also can use psychology to deter and defend.

Tamper seals

You might be forgiven for thinking that "tamper seals" are Florida sea creatures or a new baseball franchise, but they are merely a cheap method of discouraging people from messing with the internal components of computer equipment. While less drastic than a lock and no deterrent to the determined thief, a seal will nevertheless make most employees think twice. The personal computer administrator could simply place a printed adhesive label across any two parts of the case that come apart. You can print such labels on the PC itself, using a design like the one in Figure 4.12.

If the system is properly administered, a torn seal will always mean that there has been an unauthorized entry into the system unit. Persons requesting legitimate access to the innards could be issued with a fresh seal from the administrator. Of course, the seals should be serialized, bar coded, or otherwise made difficult to copy. Consider the way this idea meets the four protection criteria.

Deters: Probably. A good deterrent to the casual thief or meddler.

Delays: No. Most seals can be torn or sliced apart very easily.

Detects: Yes. Will show that security has been breached.

Responds: No. Hard to tell who broke the seal and when without heavy surveillance.

Making your mark

One of the quickest and cheapest ways to defend your personal computer hardware is to mark it. By placing a hard-to-remove serial number, identification number, or statement of ownership on your hardware, you decrease its value to the thief. Several systems are available. You can scratch or etch an identifying mark or use a special ink that is visible only in certain light. The effectiveness of a marking system is greatly enhanced by the display of prominent notices in the office, such as the one shown in Figure 4.13.




Figure 4.12 Tampering detection seal.

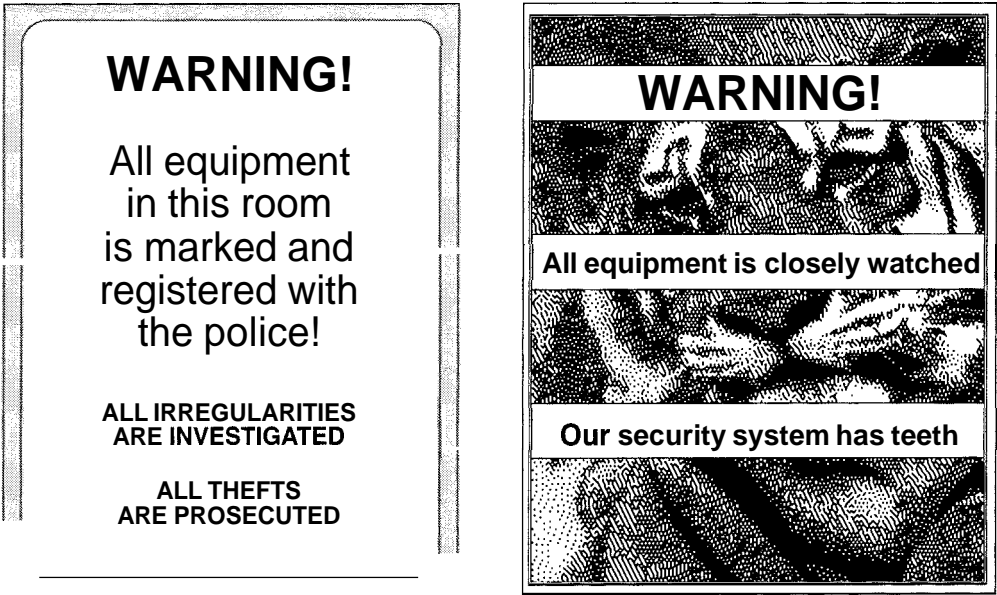


Figure 4.13 Warning signs will deter some thieves

Warning signs make would-be thieves aware of an added obstacle to fencing the goods, as well as an added means of detection should the goods be intercepted. The notice should indicate that a marking system is in use but not be specific about the system used. Of course, some of your hardware—like monitors, printers, and system units—has fairly good identification already in the form of a manufacturer's serial number. However, while you should diligently keep track of all serial numbers, you should not rely on this for your marking because some of the numbers can be removed. Furthermore, some systems, notably hand-built PC clones, might lack any form of serial number.

One effective equipment ID system is Security Tracking of Office Products, or STOP, which is operated by Missouri-based Business Security Systems. The most obvious part of the system is a metal plate that is attached to the equipment. The markings on the plate are illustrated in Figure 4.14. You then register the equipment with Business Security Systems.

Applying the plate also applies an indelible marking to the equipment, which remains even if the thief succeeds in removing the plate. Together with warning stickers like the one shown in Figure 4.13, this marking acts as further discouragement to both the theft and the fencing of equipment. The details recorded in the tracking database help identify stolen goods that are later retrieved by the police. Remember that, without proper identification records, it is often impossible to claim stolen goods, which then might end up back in the hands of the people who stole them!

You might not expect to find a software-based identification system, but Micro Law Software of Oregon has created a system called Micro ID that hides a unique identification deep within your computer's hard disk. This ID is able to survive all but the most rigorous low-level formatting (see Figure 4.15). You register the number with the authorities, who then can use a separate piece of software to identify and verify ownership of recovered computers



Figure 4.14 Equipment identification and tracking plate from Business Security Systems.

Owner Identification Information:		Level of Accuracy = High	
Name: JOHN DOE		DOB or Date Marked <03-12-1963	
28754 B AVENUE			
TROUTDALE, OR 97060		Phone #: 503-661-1211	
Computer Identification Information:			
Computer Brand Name: GATEWAY 486-66		Serial #: 18745-445673-A44	
Computer Component Information:			
The following listing reflects differences in components from when MICRO-ID was originally used on 09-10-1994 and today 03-10-1996 at 14:16:05. MICRO-ID was provided to the owner by Gresham Police Department, (503-661-2222).			
When	DOS Ver	CPU/Math Chip	Floppies/P Ports/S Ports/Monitor Information
Then	5.0	486 Yes	Two One Two VGA Card
Today	6.0	486 None	One One Two VGA Card
This analysis was made by COP-Only Version 3.0 which is owned by MICRO LAW SOFTWARE, INC. and is licensed for use by THE WASHINGTON COUNTY SHERIFF ONLY. Use by any other agency or entity is in violation of Federal Copyright Law!			

Figure 4.15 The Micro ID system from Micro Law Software

Consider the way that equipment identification stacks up in the four categories:

- **Deters:** Definitely. Makes the hardware harder to unload. A sign pointing out the use of a marking system can be preemptive.

Delays: Definitely. A thief might think twice about running off with your marked hardware as opposed to someone else's unmarked equipment.

Detects: Possibly. The authorities can better trace and identify your property if it is marked, and this might help them build a case against a perpetrator.

Responds: Yes. Enables potential buyers and the authorities to spot stolen goods and return to rightful owner.

Faking it

The greatest amount of apparent deterrence for the least amount of money might be conspicuous warning signs and messages of foreboding. A sign like the one in Figure 4.16, posted on the outside of an office door, might cause some intruders to think twice. If only a few people know that the message is not entirely true, then you can deter would-be wrongdoers without spending a lot of money. Only those on the inside need know that the company cannot afford the real thing.

The signs themselves can easily be made up using the computer itself. Now that many personal computer printers can create almost typeset quality output, signs can be fabricated very quickly. You can use clear plastic adhesive tape to fix the signs onto equipment, furniture, walls, or doors. If you use wide sheets of plastic that cover the sign without joints, then the sign appears to be laminated, and a very convincing result can be obtained. If your office actually has a laminating device, then enclosing signs in stiff plastic will add to the realism.

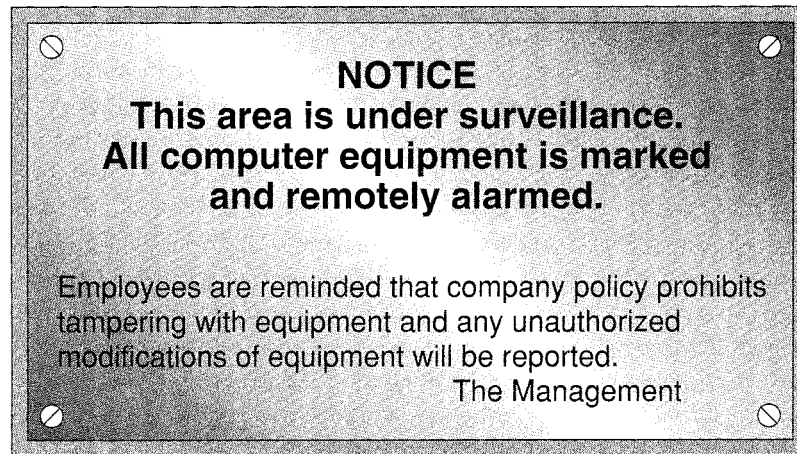


Figure 4.16 A warning sign can be effective even if it is not completely truthful

Live warnings

A variation on this theme is to use computers themselves as warning signs. A colleague who works from home was about to take a vacation. Worried about the possibility of a burglary, he planned to lock his personal computer system in a closet before he left. However, he realized that, because his personal computer also is his fax machine, he would have to leave it turned on while he was away. Faced with this dilemma, he wrote a small batch file program that displayed the screen seen in Figure 4.17.

The only way that you can test whether or not this message is true is to risk setting off the alarm. Realizing that there could be power failures while he was away, he altered the computer's AUTOEXEC.BAT so that the message was displayed whenever the PC was turned on. However, he realized that simply displaying the message




Figure 4.17 A batch file warning message.

would not be a good idea because prolonged display of the same single image in the same place might damage the monitor.

To get around this problem, several different messages were created and displayed in turn, with a short gap between them. In fact, the effect of screen movement added to the realism of the warning. This was further augmented by using beeps from the computer's speaker. The final result was an impressive message that probably would discourage all but the most sophisticated burglar. See the notes in appendix B for more about building batch files like this. You will find this example and several others on the Cobb/NCSA Security Resource disk included with this book.

Warning Windows

If you use Microsoft Windows, you have a built-in alarm warning in the Marquee screen-saver module. This is activated through the Desktop module in the Control Panel, which usually is located in the Main program group. This screen saver blanks the screen after a set period of time, then continuously scrolls a piece of text across the screen. By using the `setup` button you can change the size, font, and color of the text as well as the color of the background. You also can enter your own text, such as `Alarm On, Motion Detector Armed!`

In fact, this is not complete exaggeration. Any movement of the mouse or any key on the keyboard will prompt the screen saver to restore the screen, thus offering a primitive form of motion detection. You can create quite an impressive effect by designing a warning sign as wallpaper and minimizing all program windows, as shown in Figure 4.18. When the screen saver restores the screen, the first thing the interloper sees is further evidence of a security system.




Figure 4.18 Desktop warning in Microsoft Windows

Wallpaper is simply a bit-mapped paint file in BMP format selected through the Wallpaper setting in the same Desktop module that controls the screen saver. I created the one in Figure 4.18 with the Paintbrush program that comes with Windows (a copy is on the Cobb/NCSA Security Resource disk that comes with this book). After you load Paintbrush, use the Options menu to set the Image Attributes so that Width and Height, measured in "pels," match your current screen resolution setting (as in 640 x 480, 800 x 600, and so on). Now, fill in a background color and create your text (you will need a large font size).

Save the file as WARNING.BMP in the Windows directory, then use the Desktop module in the Control Panel program to select this file as the Wallpaper setting. Check the Center option, and you are all set. Minimize any programs that you are running to make the warning easy to read. Leave your desktop like this, and when the screen saver kicks in, your "alarm" will be in place. Moving the mouse or pressing any key will instantly reveal the desktop warning (this is particularly effective in a dimly lit room if you have used a dark background for the screen saver and bright background color for the wallpaper).

Securing the Right Thing

One of the first computer-related crimes that I encountered was a nighttime break-in at the office of a planning consultant. This particular case illustrates the need to consider the security of the personal computer system as a whole and the need to think carefully about what needs protecting.

The consulting business used two IBM PC-XTs. One of these was on the desk of the receptionist/administrative assistant who did most of the typing of the lengthy planning documents that the consultant prepared for his clients. This desk was visible from the path that led to the car park that served the building in which the office was located. This probably was where the thief got the idea for the break-in. An alarm on the premises did go off when the burglars entered, which probably accounted for the fact that only one PC was taken and that it was very hastily ripped off the desk with some cables still attached. However, the police response was not quick enough to be of much use in apprehending the criminals.

When the consultant arrived at the office to survey the damage, he was relieved to find the set of backup data disks in his desk had not been taken. He arranged for a rented PC XT to be delivered to replace the stolen one. The large letter-quality printer, an IBM 5218, still was in place. When his assistant arrived, they thought that they would be back in business very quickly. However, when they installed the rented computer, they realized that the printer cable was missing. It only took a few phone calls to discover that this was no off-the-shelf cable. As the consultant tried to contact IBM for a cable that would attach an IBM PC XT to an IBM Model 5218 printer, his assistant installed their word processing software, DisplayWrite 3, on the rented XT. A bid submission deadline was fast approaching, and the office had no printing capability.

The consultant called around to rent a printer and found himself sinking into a nightmare. Very few printers were compatible with DisplayWrite 3; indeed, none of the units on offer from rental companies would work. At that time, DisplayWrite 4 had just been released, and it supported more printers. DisplayWrite 4 reads Dis-

playWrite 3 files. They could upgrade to DisplayWrite 4 and use it with a rented printer, but upgrades took weeks to process. They could go out and buy DisplayWrite 4, but there would be a learning curve to get used to the program changes. Besides, DisplayWrite 4 did not support the IBM 5218 printer that they wanted to use as soon as a cable was found.

Finally a man from IBM arrived. He installed and tested a cable. It did not work. A new serial card was installed. A second cable was tried. It did better. A configuration file used to redirect print output to the serial port had to be fine-tuned. Finally, the printer communicated with the PC. However, DisplayWrite 3 still would not print pages properly. Another cable was tested. Finally, it was determined that IBM had issued a special memory-resident driver program for the 5218 printer. This was needed because the 5218 printer had been created for the original DisplayWriter dedicated word processor, which did not use DOS. A copy of that driver program had been installed on the stolen PC. After a frantic search, the consultant found the master copy of the driver. However, simply installing it did not solve the problem. To run correctly, the driver program needed certain settings, and the documentation was nowhere to be found.

Eventually, several calls to IBM later, the correct settings were determined and the problem was resolved. Yet this was two agonizing, frustrating, and expensive days for the consulting company. What lessons were learned?

Backup data files are not enough

Some of your backups must include installed program files. The vital driver software and configuration file that let the printer and computer understand each other took hours to fine-tune by trial and error but were obviously already installed correctly on the missing computer.

One-off equipment is invaluable

The stolen cable was a custom job, very hard to replace, and valuable far beyond its dollar cost. Without it, the office lost its printing capability. The harder something is to replace, the more closely it should be protected. The other side to this is to avoid specialized hardware. Clearly this goal cannot always be achieved, but there are definite advantages to using readily replaceable components.

Don't overlook commonsense security

In this example, the placement of personal computers near the front door probably made them a tempting target. Thieves have plenty of targets to choose from, so make yours less appealing. The alarm system did help. Only one computer was taken. Even that might have been left behind if it had been further from the front door or attached to the desk.

Key Management

Many physical security devices use removable keys, as do some system access controls, such as keyboard locks. These keys must be properly managed. A system of

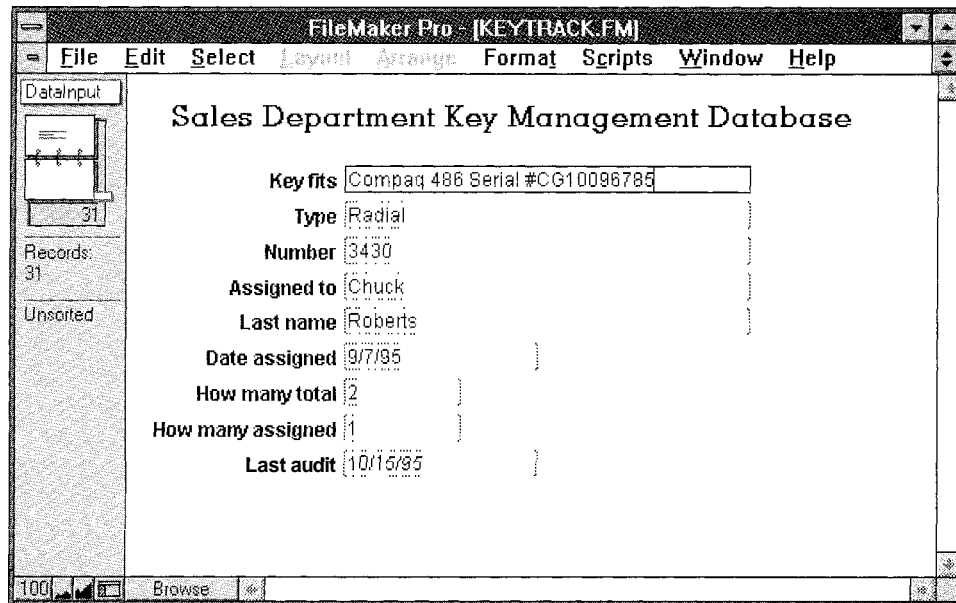


Figure 4.19 A database for tracking keys.

key management will ensure that keys are not lost, improperly distributed or copied, or retained by departing personnel.

Keeping the keys

Any serious security program that relies upon physical keys must include some form of key management system. There is nothing exotic about key management; it is simply one of those tedious record-keeping systems that works only if you work at it. You might find that a personal computer database can help keep track of key assignments. In Figure 4.19, you can see an example of such a database, which was created with FileMaker Pro.

The main elements of a key management system are as follows, with some applying specifically to keyed equipment, while others apply to keys in general:

- Separate keys from locks as soon as possible. Make sure that all of the keys that come with a piece of equipment are controlled.
- Cross-reference keys to locks using a protected code. Do not label keys so that they can be matched to locks without the code.

Only issue keys as they are needed. Do not leave keys in locks.

- Collect all of the keys from employees who are leaving before they leave.
- Be prepared to change locks if keys are missing.

Be alert to the fact that keys can be duplicated. Mark keys "do not duplicate without permission."

Establish a relationship with any locksmiths who have premises near your site to minimize the chance that they will perform inappropriate duplications.

- Mark internal keys, such as those for keyboards, with the name of your organization. This will help tracking and retrieval and alert locksmiths. Do not mark the front door keys in this manner.

Reward employees for looking after keys. Penalize them for negligence. Let employees know that they will be held responsible for fraud/abuse carried out on a system for which they have a key.

Remember that possession of a key does not guarantee the identity of the key holder. See chapter 6 for more on identification and authentication systems.

Key management terms and technology

Don't be surprised if some of the previous points strike you as mundane and low-tech. A recent issue of the magazine *Security* observed that "At a time when automation and accountability are standard buzzwords for many areas of security, key control has remained, by and large, a manual endeavor." However, that is beginning to change, with computerized key safes and drawers beginning to enter the market. For example, the Keywatch System from Connecticut-based Morse Watchman offers a "smart" storage cabinet that requires a PIN number to access a key. This helps to track who has a particular key and when they have it.



PIN Money

Note that the term PIN stands for personal *identification* number, so "PIN number" is actually a redundant expression. However, it is widely used, particularly now that most people who have a bank account also have a cash card that can be used, in combination with a PIN, to withdraw money from an automatic teller machine (ATM).

From Key Systems of New York comes Security Key Monitor, which is a locked panel where keys are released only to authorized users. An audit trail keeps track of who has which keys. It is possible to assign time limits to keys, and an alarm will notify security personnel if a key is out too long. Like the Keywatch system, Security Key Monitor can handle from a dozen keys to several hundred.

If you have from 100 to 3600 keys, then Key-Trak of Florida offers a fully computerized system of key tracking, which is operated by a PC. The Ford dealership where I purchased my last car uses this system to track hundreds of new- and used-car keys. The keys are placed in a locked drawer that is computer-controlled much like a cash register. When an authorized user opens the drawer, the system records any keys removed as "checked out." Authorized access to the system can be based on passwords, magnetic ID cards, bar codes, or proximity readers, allowing integration into other access control systems (as discussed in chapter 6).

When you are dealing with keys and locks, it helps to know some of the terms that are commonly used, particularly when ordering a number of locking devices:

Keyed differently: All of the locks in a set of locks require different keys. In other words, each lock requires its own key.

Keyed alike: All of the locks in a set of locks use the same key. In other words, one key opens all of the locks.

Master keyed: All of the locks in a set of locks require different keys, but there is one special key that will open all of the locks.

Reserved keyway: The lock is designed so that the keys to it cannot be duplicated. In other words, the key is not based on any of the commercially available key blanks.

Computer Insurance

Anyone who has had to file a claim on an insurance policy will tell you that insurance should not be your first line of defense. Claiming on an insurance policy is a chore even when dealing with the best of insurers. Furthermore, many insurers will want to know that you have taken risk management seriously before they will offer full coverage. Insurance is a complex subject, and this book can give you only general guidance on insuring the computing aspect of your personal or business activities. If you have an insurance broker that you trust, he or she is the best person to advise you on the right policy for your particular circumstances.

Determining coverage

The first question to ask is: "What insurance do I already have?" If you have a computer in your home, it might be covered by your home contents insurance. If we are talking about computers in an office, then there might be some coverage already in place. However, you need to be sure about what is covered, against what threats it is insured, and what exactly the insurance policy will pay if any of those threats materialize.

Hardware. The first item on the list usually is hardware. This is relatively easy to identify, record, and value. You should insure it for replacement cost so that, if it is stolen, for example, you can immediately, and at no loss to you, go out and buy something to replace it, thus minimizing losses from lack of system availability.

The range of threats for which insurance can be obtained includes theft, accidental or intentional damage, water damage, fire, power surges, and lightning. To obtain insurance on your hardware, you might have to describe it in detail and install suitable precautions, such as smoke detectors.

If you already have some form of office or home contents insurance, do not assume that it covers all of your hardware for all of these threats to the desired value. Some home policies exclude equipment used for business, and they might limit the "per item" insured value of big ticket items. You might have to declare and list computer equipment for it to be covered, or you might have to purchase a "rider." Thus is a supplement to the original policy for a specific item or set of circumstances.

Avoid the temptation to purchase a policy without being entirely clear about what it covers. When it comes to existing policies, do not tempt fate by "hoping for the best." You are likely to get a nasty shock when it comes time to claim. In other words, do not assume anything, and buy your insurance on the premise that, at some point, you will have to rely upon it. Be sure to ask the insurer or the agent how claims are handled, how losses are valued, and how payments on claims can be applied. (For example, if an aging 386 is stolen do you have to replace it with an aging 386 or can you apply the replacement value of that machine against a newer model?)

Software. For insurance purposes, software may be divided into purchased and created. For example, when a policy says it covers "purchased software," it means that, if someone steals all of your copies of dBase, you can get replacements for as many licensed copies of dBase as you purchased. It does not mean that you will get paid for 10 dBase licenses if what was stolen amounted to one licensed and nine unlicensed copies. The term "purchased software" also excludes the applications that you developed in-house with dBase. Custom applications can be insured but only if you declare them and value them.

The threats to software parallel those to hardware, with perhaps one addition: viruses. It is possible to purchase insurance that covers loss of software from a virus infection.

Data. The data that is stored on a computer rarely has a clearly quantifiable street value, and insuring it accurately is going to be difficult. The typical approach is to estimate the "person hours" required to recreate the data. This can be done at several levels, for example, based on losing all copies versus losing all of the changes since the last backup, if that survives. Remember that data is not likely to be covered as software, and do not assume that data is covered by your policy unless there is specific language, or a rider, to that effect.

Insurance tips

If the previous discussion strikes you as a rather grim picture, it is nevertheless realistic. However, there are some bright spots. In some circumstances, your personal computers might already be covered against certain threats. For example, if you have an automobile policy that covers "theft of items from inside a locked vehicle," it might cover loss of computers that are in transit.

Notebook computers present a particularly interesting case as they are more prone to travel and be damaged by accidents, such as uncontrolled descent in the direction of unyielding surfaces. Fortunately there now are special policies for portable computers, even when they are being taken overseas.

Most insurance policies are likely to have a deductible, an amount you have to pay towards a loss before the insurer will pick up the balance. If you are prepared to accept a higher deductible, you can probably get a lower premium.

Be sure to include peripherals, like modems, scanners, network adapters, mice, and cabling. Be sure to keep the insurance company informed of any changes, such as upgrades to storage and memory capacity. When disaster does strike, you don't want to discover that new acquisitions are underinsured.

Check what the policy says about reinstatement of damaged equipment. Typically, goods are replaced on the basis of "similar equipment in a condition equal to but not better or more extensive than its condition when new." With the rapid rate of change in computer technology, you might be claiming for equipment that is somewhat obsolete or even out of production. This means that the replacement model is likely to be better, and the policy might require you to pay all or part of this difference. For larger policies, this point might be negotiable.

Decide whether you need to insure for "business interruption." This usually is regarded as having two distinct elements. The first element is the "increased cost of working," which means the cost of getting through the usual workload without the computer, which would include the cost of renting new equipment and hiring temporary staff. The second element is the "consequential loss" (for example, the loss of revenue because files cannot be accessed and so invoices cannot be sent to customers). Business interruption is another area where quantification is difficult, and professional advice geared to your specific circumstances is essential.

Finally, check what the policy says about "reasonable precautions" to prevent or minimize loss. These precautions should be genuinely reasonable. They might include due care in selection and training of staff, compliance with hardware maintenance agreements, and normal antitheft precautions. Make sure that the provisions are not so extensive as to provide loopholes for an insurer looking to avoid paying a claim. For example, an insurer might refuse to pay the full cost of recreating data if a hard disk is damaged by a power surge and a proper backup regimen was not followed. If the entire office is destroyed by fire, including the backups, you might think that the full replacement cost of the data is a reasonable claim. However, the policy might specify that backups be stored in a fireproof safe. In general, following a good risk-management plan and sound security policies is both a prerequisite to reasonable claims settlement, as well as the best way of avoiding the need to claim.

Costs and sources

If you find that your current insurance does not adequately cover your personal computer resources, you can purchase special computer-related policies. Apart from calling the broker that you use for your other insurance needs, you can purchase computer insurance by telephone, direct from companies that specialize in this field.

Determining the likely cost of insurance is rather like asking how long is a piece of string. A lot will depend on the value of the hardware insured, along with the other risks insured on the policy, particularly business interruption. A further factor will be the individual circumstances of the policyholder; users with bad claims records or businesses located in high-risk areas might pay more.

As a guideline, a generic direct policy that offers replacement cost coverage of hardware and off-the-shelf software against fire, theft, power surge, lightning, flood, and earthquake damage will cost around \$100 per \$10,000 of insured value per year, with a \$50 deductible. For larger coverage, you might try using an independent insurance broker or intermediary to find the best quotes for you. However, direct computer insurance specialist SAFEWARE also offers high-end coverage that takes in items such as data re-entry.

If you use an agent, he or she should be prepared to "rebroke" for better quotes when the policy comes up for renewal. Make agents earn their commission by handling any claims that you subsequently have to make. They know the shortcuts through the insurance industry red tape and can speed things up considerably.

When taking portable computers overseas on business or even personal trips, you probably are removing them from normal coverage. You can purchase international coverage from SAFEWARE on a 90-day or annual basis. The rate currently is \$22.50 per \$1000 coverage for 90 days and \$50 per \$1000 for annual coverage. The deductible is \$250, and most threats are covered, with the notable exception of damages relating to power supply (perhaps there are too many people who, like me, at some time, have plugged a 110-volt device into a 220-volt supply).



SAFEWARE Then and Now

In 1990, I obtained the rate of \$100 per \$10,000 from SAFEWARE, an Ohio-based company specializing in computer coverage. When I was revising this book in 1995, I called SAFEWARE for an update. The rates have stayed exactly the same, and David Johnston, the company's founder and chairman, anticipates that they will remain unchanged for the next five years.

What has changed in the meantime is my perception of "direct" insurance offerings. I now feel much more comfortable suggesting SAFEWARE to people who need computer insurance. The company has managed to hold the line on the cost of coverage, while increasing the range of its offerings and maintaining the staff levels required to service claims. Let's put it this way: Yes, my computers are insured by SAFEWARE, and no, I don't get a special deal for saying that.

Removable Media

To a lot of people, the topic of information security still is synonymous with hackers, viruses, tiger teams, and technology that sometimes is more fiction than science (a "tiger team" refers to people employed to test security measures, as depicted by Robert Redford and his associates in the movie *Sneakers*). However, information security typically is more mundane than this.

The flighty floppy

One of the biggest threats to the security of computer-based information is the humble floppy disk or diskette. Cast your eyes around just about any office, and you will be able to spot at least a dozen mislabeled or unlabeled floppy disks. Seldom do you see anything on a diskette label that indicates ownership or the value of the documents contained on the disk. Often there is no indication that the diskette has been approved, classified, or even virus-scanned.

To the security professional, each one of these "under-labeled" disks has LEAK and THREAT written all over it. One 3.5" disk could be 1000 business letters stolen by a disgruntled employee or perhaps 20,000 customer records sneaking down the street to your competitor's office. (There's more than one reason that floppy disks

are called SneakerNet.) All too often, there are no controls in place to prevent trade secrets, proprietary processes, business plans, and even sensitive personnel information from being copied to a diskette. With staff reductions a part of office life around the world, floppy disks can be just too tempting for some "rationalized" or "downsized" employees to resist.



Floppy Follies

- By the time American super-spy Aldrich Ames was arrested, he had taken home dozens of floppy disks worth of sensitive data.
- In 1994, American Airlines sued Northwest for the alleged theft of confidential data, mailed on a floppy disk to Northwest's head office in Minneapolis by an American Airlines employee, who went to work for Northwest shortly thereafter. According to American Airlines, the data is worth \$50 million.
- The pharmaceutical giant Marion Merrell Dow sued two ex-employees for unauthorized removal of computer files containing secret processes for manufacturing Cardizem, a very successful medicine. The company described the cost of losing its proprietary technology "almost inestimable, reaching hundreds of millions of dollars."

Most data security surveys indicate that more than half of losses are due to insiders, and floppies are an ideal medium for clandestine data transportation. The very qualities that make them so useful for legitimate data transfers, such as size and convenience, also add to their appeal for less innocent purposes. You don't have to be a computer security expert to understand why floppy disks are so suitable for larcenous purposes. Ask any business how it handles the cash it keeps on hand. The most likely answer will be "very carefully."

Hundreds of years of accumulated experience dictate that we track cash diligently and, whenever possible, convert notes and coins into safer assets, such as a bank balance or certificate of deposit. Unfortunately, floppy disks have been around for only a couple of decades, and although many organizations now have realized that information is a valuable asset, few appear to have taken to heart this simple equation: "Floppy disks are to information as cash is to money."

Just how much information are we talking about? Consider this number: five billion megabytes! That is the total storage capacity of the four billion 3.5" floppy disks sold in 1994. That's five million gigabytes of data storage, most of which, according to what we have seen of today's standard office procedures, will not be properly labeled, tracked, or audited. In other words, there has to be a major change in the way that management and employees regard floppy disks.

D-Day

While a slow and steady piecemeal approach might be less threatening to employees and more appealing to management, one of the best ways to raise floppy disk consciousness is an instant audit to identify and catalogue each and every floppy disk on the premises. This requires careful advanced planning so that, when you announce

"Today is Disk Day" and request department heads to log all disks, they have the necessary tools to do the job.

Some form of consistent disk labeling should be introduced at this point so that all disks that have been logged and virus-scanned are clearly identifiable. Disk Day also is when you introduce a new regime of disk control whereby all incoming disks are checked by appointed personnel. All other disks start out freshly formatted (so that there are no hidden or deleted files) and are signed out to individuals who then are held accountable for them. Virus scanning should be an integral part of this process so that no disks are authorized without being scanned. These controls, plus the mandatory labeling of disks, will ensure that unauthorized disks are easily spotted and immediately suspect.

When it comes to disk labeling, some companies already use their own inhouse labels, but this concept can be taken several steps further. An American company, Avanti Associates, has devised an extended label system, called TAG, that allows you to put not only the company name on "approved" disks, but also a data classification and warning statement, plus a serial number and bar code unique to each disk. At the bottom of Figure 4.20, you can see a preprinted label ready to be removed from its adhesive backing.

The first two sections on the left are folded over to form a tag, as seen at the top of Figure 4.20, which shows the label attached to a disk. The section on the right is the part that actually attaches to the disk. When the label is removed from the backing, it leaves behind a matching bar code sticker that can be placed on a log sheet so that you can record to whom the disk was issued.




Figure 4.20 The TAG label attached to a disk and ready to be removed from its adhesive backing.

Read the label

Past efforts at creating more meaningful labeling have been thwarted by the physical limitations of disk size. Avanti has come up with a simple but effective solution to this problem. The extended label tags actually hang out of the disk drive when the disk is in use. This has the added advantages of increasing the visibility of disks and helping to prevent floppy disk booting, which is a major cause of new virus infections. (According to the UK publication, *Virus Bulletin*, for the past three years the majority of reported virus infections have been caused by boot sector viruses.)

Virus prevention is a major benefit of any well-designed removable media control system because floppy disks are the main carrier of viruses. A survey by the NCSA indicated that more than 75% of virus infections come from floppy disks, and most of these are disks entering the workplace without being scanned or authorized. Thus, a disk control regime will not only help to reduce data leakage, it will substantially decrease the chances of virus infections.

However, the primary benefit of improved disk labeling and disk authorization is a widespread increase in security awareness and compliance at the user level. If management is seen to put a value on disks and pays serious attention to tracking them, this will have a knock-on effect to all members of the organization and any visitors who bring disks with them. (There is one well-known financial institution in London that will not allow any disks to enter the building unless they have been scanned for viruses, even disks that clients happen to be carrying with them.)

A slightly more high-tech system, which could be implemented in concert with TAG, is supplied by Reflex Magnetics. Known as the Disknet system, this is a multi-layered security solution. Once the Disknet system is installed, users are allowed to work with only special floppy disks that are electronically encoded during manufacture. The Disknet software monitors any floppy disks that users put in their machines and denies access to or from the disk if it is not approved.

At this point, you might be wondering "How would better labeling of the disks, or even disk serial numbering, actually prevent someone taking a disk out of the office?" The high-tech answer is to use the TAG system with radio-sensitive inserts, similar to those used to prevent shoplifting or to control library books. Scanners at exit points set off an alarm if disks pass through them. The low-tech answer is that, by raising awareness of the crucial importance of information to the organization, and thus the future of jobs, and reinforcing this importance through security policies that are both declared and enforced, you will create a climate in which users think twice before transgressing. Your efforts can be underlined by making security compliance a part of employee evaluations and by taking disciplinary action against those who do not comply.

Disk authorization schemes have many benefits, including the prevention of software piracy, which now is a major cause of exposure for companies. Six-figure fines are not unusual when an organization is found to be using unlicensed software. Controlling the movement of floppy disks in and out of the office is an essential part of any program to reduce piracy. For example, you might dictate that employees who want to bring disks into work must submit them to license checking as well as virus scanning. That can help to prevent one application propagating throughout the building when no more than one license actually has been purchased.

One further benefit of a properly implemented disk control system involves the legal term "standard of due care." In other words, if practical and economical means exist to prevent the loss of data, they should be employed. If employees are hired away by a competitor and take disks full of data with them, a key factor in the ensuing lawsuit is likely to be the measures employed to prevent such occurrences. If you don't mandate that confidential information should always be marked "Confidential," it will be a lot harder to argue in court that the information actually is confidential.

Systems for tracking, authorizing, and monitoring floppy disks within an organization exist, but it is fairly obvious that current levels of deployment are minimal. Is this because they are too expensive or too much work? Or is it simply because we have not yet learned the true value of tracking and controlling removable media? People who have suffered data leakage through floppy disks will gladly tell you that spending substantially less than a dollar per disk to prevent losses is far better than having to file expensive legal actions or waste thousands of person-hours recovering what you have lost. Hopefully, not everyone will have to learn the hard way. By stressing the advantages of removable media control systems, as well as publicizing what can happen in their absence, we should be able to achieve a significant increase in the general level of floppy disk awareness.

The Network Connection


When you network a bunch of personal computers to share storage facilities on a machine that operates as a file server, physical protection of that machine becomes especially important. File servers are a juicy target for today's increasingly computer-literate crooks who know that file servers tend to contain more RAM, faster processors, and bigger drives than ordinary desktop machines. *All* too often, we see file servers sitting on or under a desk in an insecure area. Sometimes they are tucked into a corner getting dusty.

You have to secure your file server, preferably in a locked or restricted access room, not just as a precaution against theft but also as a defense against abuse from internal hackers, malcontents, or misadventurous supervisor-wannabes. In chapters 11 and 12, you will learn that many attacks on networks can be prevented if physical access to the server is protected. Locking it away and chaining it down is a small price to pay for the continued confidentiality, integrity, and availability of networked data.

Summary

Physical security is the most basic level of security. This is important to bear in mind as we get further into data and program security. Sophisticated file and system access controls can ensure confidentiality when media or computers are stolen, but they offer little defense against the threat to availability that outright theft presents. Vandals who break into your office and trash your hardware are not deterred by encryption schemes. Overall office security is the first line of defense. The safety of an organization's personal computer equipment and the data stored therein depends upon the organization taking appropriate steps to secure the premises. Beyond this lies protection from the enemy within and methods to minimize the impact of breaches in site security.


Dilbert by Scott Adams, reprinted 4 of
United Feature Syndicate, Inc. rmis



S. Adams E-mail: SCOTTADAMS@AOL.COM



5/19 © 1995 United Feature Syndicate, Inc. (N



Secure Power Spikes, Sparks, and Electrical Threats

W.E.B. connection: <http://www.ncsa.com/pclan/chap05.html>

*“ they want rain without thunder and
lightning They want the ocean without the
awful roar of its many waters ”*
FREDERICK DOUGLAS

This chapter addresses threats and risks that arise from the fact that computer hardware runs on electricity. Our first concern is ensuring a smooth and constant supply of electricity. This is not just something for the office electrician to think about. Every personal computer user needs to know about surges, spikes, and overloads. If the supply of electricity to your hardware is interrupted unexpectedly or varies significantly from expected parameters, the consequences can be serious. There are several other security implications that arise from the fact that computers run on electricity. These include radio interference, which can be used for eavesdropping and sabotage, plus radiation, which is a potential liability threat.

Power to the Computer

The consequences of abnormalities in power supply include loss or corruption of data and permanent damage to hardware. As you can see from the chart in Figure 5.1, losses attributable to surges run into the hundreds of millions of dollars.

These figures were supplied by SAFEWARE, the leading source of computer insurance in the U.S. Each year, SAFEWARE publishes its estimates of total losses based on the experience of insuring approximately one out of every one thousand personal computers in America (note that the \$318 million worth of losses due to surges does not include an additional \$86 million in losses attributable to lightning).




Figure 5.1 Chart of losses by category in 1993 (Figures provided by SAFEWARE)



Lightning Statistics

At any given instant, there are more than 2000 thunderstorms taking place throughout the world. All of these storms combine to produce about 100 lightning flashes per second, each one with a strength of up to a billion volts and temperatures of more than 54,000 degrees Fahrenheit. At the height of a moderate-sized thunderstorm, it can generate several hundred megawatts of electrical power, or the output of a small nuclear power plant.

NASA, Kennedy Space Center Lightning Research Program

Electricity 101

Fortunately, there are devices that can help you protect against these dire consequences. They go by names like *surge suppressor*, *line conditioner*, and *uninterruptible power supply*. Before examining how these devices can help, it is important to be clear on some of the basics of electricity. Personal computers consist of a large number of electrical circuits. Some circuits, such as those that hold information in memory, use only a small amount of power. Other circuits, such as those that power the disk drives, require more electricity. The electricity used by the internal components of the personal computer is known as *direct current* (or dc) because it is always traveling in the same direction (the wire supplying the electricity is positive, the return wire is negative).

As a rule, personal computers get their electricity from normal domestic electrical circuits, sometimes referred to as *mains current*. This mains current is fairly strong and is known as *alternating current* (or ac) because it alternates between positive and negative. Within most personal computers is a piece of equipment called a *power supply*. This takes in the ac current from the mains and converts or trans-

forms it to low-powered dc used by the computer components. You can see this diagrammed in Figure 5.2.

The power supply is a vital component in any personal computer system, and the one that bears the brunt of abnormalities in the mains supply. (The term power supply will be reserved for this component and will not be used for other devices, discussed later, which also can be said to supply power.) A notable and growing exception to these observations are notebook computers, which run on batteries. Notebooks are discussed later in this chapter.

Electricity normally is useful only when it is allowed to flow through a device or appliance. For example, if it flows through a lamp, it produces light. If it flows through a disk drive motor, it produces rotation. Regardless of what it flows through, electricity produces heat, in quantities that vary according to the appliance. For example, when electricity flows through a normal household bulb, the bulb not only gives out light, it also gets hot. An electric blanket is an appliance in which electricity does nothing besides produce heat. The electrical circuits inside a personal computer generate heat, which is why many systems are fitted with fans that provide air flow for cooling.

Measuring the power

The reason that electricity flows through an appliance is a difference in electrical pressure between the wire supplying the electricity, the live wire, and the wire taking it away, the neutral wire. The greater the difference in pressure, the greater the flow. However, the flow also is determined by the electrical impedance of whatever lies between the live and neutral wires. If this impedance is low, a given pressure difference will result in a much higher flow than if the impedance is high. The term resistance sometimes is used instead of impedance, although there is a technical difference between the two. If the live wire is carrying a lot of electricity to an appliance and the appliance uses that electricity, it creates a resistance.

The flow of electricity usually is called current and can be measured in amperes. The abbreviation amps usually is used and is designated simply as A. The more amps an appliance uses, the more electricity it consumes. You will find the amperage of




Figure 5.2 The power supply arrangement in a typical PC.

many appliances listed close to where the power cord enters the casing. For example, on an Apple LaserWriter IINT, it says 7.5 amps. The pressure difference in a circuit is called voltage because it is measured in volts. The higher the volts, the more pressure there is to make the electricity flow. Volts often are designated as V, and large quantities are measured in kilovolts (kV) so that 2000 V equals 2 kV. Impedance, or resistance to flow, is measured in ohms, represented by the symbol Ω . The relationship between the three quantities can be expressed mathematically: Resistance is equal to the pressure difference divided by the rate of current:

$$\text{ohms} = \frac{\text{volts}}{\text{amps}}$$

This means that, if you know two of the amounts involved, you can work out the third.

In measuring electrical devices such as computers, a fourth quantity often enters into things: power. This can be measured in watts (designated W). For example, the bulb in a desk lamp might be 60 watts or 60 W. For larger appliances the watts are measured in thousands, or kilowatts, designated kW. You can see this and other terms diagrammed in Figure 5.3.

The power or wattage of an electrical appliance is a measure of how much electrical energy it consumes in a given period of time. For example, the original floppy disk IBM PC systems had power supplies rated at about 50 W, less than many light bulbs. This was not sufficient to run the early hard disks, which used a lot of power. Most of today's desktop machines have at least a 200-W power supply. Network file servers use even more powerful units. For most practical purposes, wattage is equal to the voltage (volts) of the supply multiplied by the current that the appliance draws from that supply (amps). The equation looks like this:

$$\text{watts} = \text{volts} \times \text{amps}$$

Suppose you have a laser printer rated at 7.5 amps. Operating at 120 volts this appliance will draw 900 watts. To turn the equation around, amps are what you get by dividing volts into watts:

$$\text{amps} = \frac{\text{watts}}{\text{volts}}$$

If your personal computer has a 120-watt power supply and you are running on 120 volts, then you might put the amperage at 1. However, this is not always an accurate equation, because the appliance might not consume the entire current that is fed to it.

For this reason, a further measure is used in discussions of computer power supplies: volt amps or VA. Like watts, this measurement also can be defined as volts times amps, but it takes into account something called the powerfactor. An appliance like an electric heater represents what is termed a purely resistive Load—one that consumes all of the power supplied to the appliance. However, a power supply in a computer system appliance is partly capacitive and partly inductive. This means that the computer system appliance doesn't use all of the power that it gets; it side-

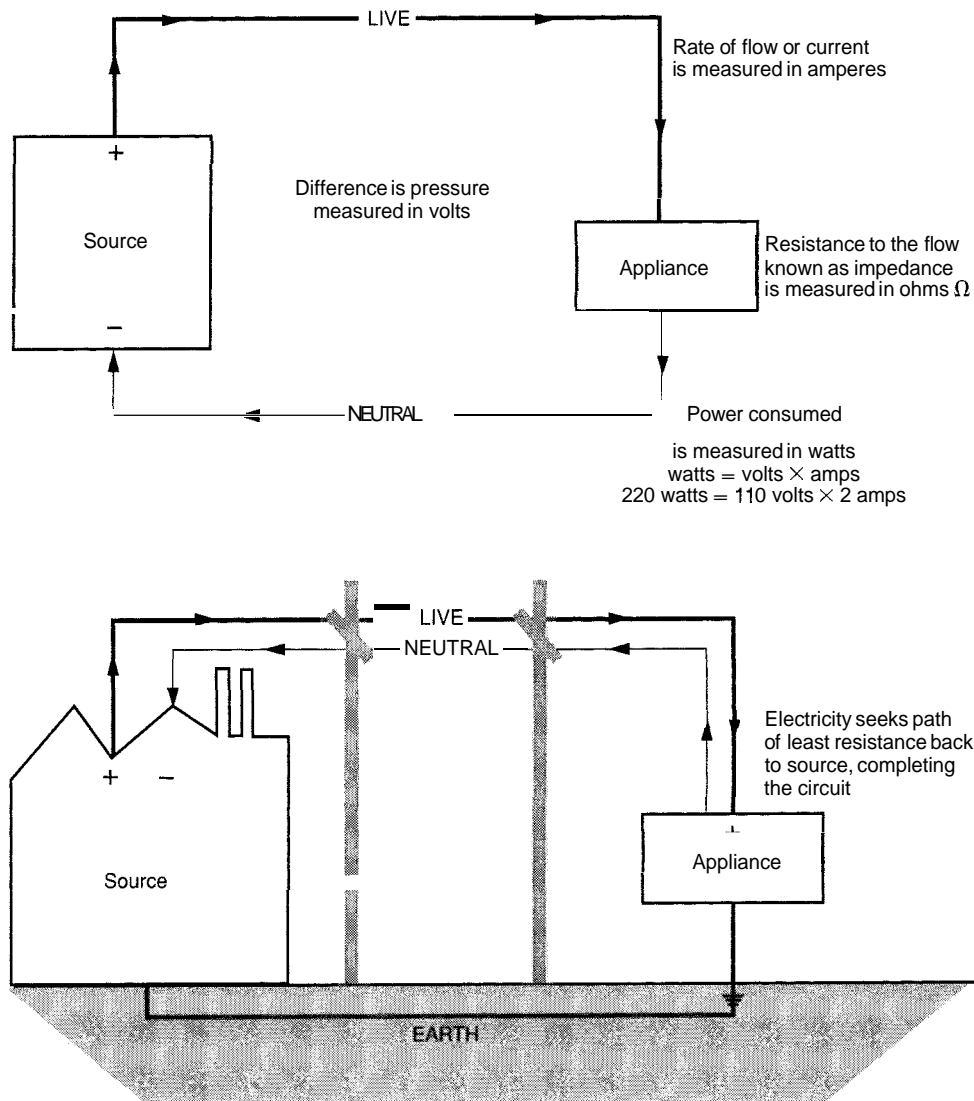


Figure 5.3 Diagram of electrical terms

lines some of it and stores some of it. This gives rise to a power factor for appliances that are not completely resistive, based on their relationship to a power requirement of 1, which represents a purely resistive appliance. Typical personal computer power supplies have a power factor of around 0.65. The equation for volt amps is:

$$\text{volt amps} = \text{volts} \times \text{amps} \times \text{power factor}$$

When amps are not known, then you can use this equation:

$$\text{volt amps} = \text{watts} \times \text{power factor}$$

Suppose your personal computer's power supply is rated at 200 watts and is operating on 120 volts. With a power factor of 0.65 the volt amps would be 130.

A matter of time

A unit of measurement that you might see mentioned in product literature for power conditioning products is joules. A joule is a fairly complex measure of energy. In electrical circuits, you can assess the energy in a flow of current for a period of time. This sometimes is useful when comparing sudden rushes of power across a circuit.

The time during which an electrical current flows can be of considerable significance. Your electricity bill probably shows the amount of electricity that you have consumed since the last bill, measured in thousands of watts per hour or kilowatt hours (kWh). If you leave 10 lamps on, each burning a 100-W bulb, for 10 hours, they will consume 10 kWh.

Electrical activity within a computer takes place at very high speed, so it is often measured in parts of a second. For example, one-thousandth of a second is one millisecond (ms). The time that it takes for a hard disk drive to read a piece of information from disk into memory often is measured in milliseconds, as in "9 ms access time." A millionth of a second is called a microsecond, and a thousandth of a millionth of a second is called a nanosecond. These terms sometimes are used in describing the performance of such devices as surge suppressors and standby power supplies.

A particularly important area of time measurement in computers involves the time taken to complete a cycle of events. Periodic time, or frequency, can be measured in cycles per second, or hertz (Hz). If you measure in millions of times per second, you use megahertz (MHz). The relative speed of the central processing chip of your personal computer is measured in megahertz. This actually is the number of "ticks" generated in one second by the clock crystal that controls the chip. The chip processes a single instruction each time the clock ticks. For example, the original IBM PC had a clock speed of 4.77 MHz. Desktop machines based on the Intel Pentium can have a clock speed of 100 MHz or more.

When the clock ticks, the power level in the clock goes from zero to full. In most desktop machines, full is 5 volts; however, in notebooks and newer desktop models, full might be 3 volts. Theoretically, this switching between zero and full happens instantly, but the voltage actually takes time to build up. The more ticks per second, the less time there is for the power to build up. To overcome this problem, the circuit is designed to apply greater pressure to get the clock voltage to rise faster. Thus, an interesting phenomenon can occur as clock speeds for computer chips are increased. The increased pressure can result in the voltage overshooting the level actually needed, placing a strain on the chips in the system. This stress, together with other design problems associated with high-speed processors, might have practical reliability/security implications.

One approach to dealing with this problem is to use lower voltages, hence the 3-volt systems. Some of these originally were designed for notebook computers where low voltages help extend battery life, but they also can be used in desktop systems, provided that you have a suitable motherboard design. They have the added benefit of producing less heat. When selecting a personal computer system that has to support mission-critical applications, there might be reliability advantages in choosing a lower voltage chip.



Harmonious Chips?

The problems of high-speed chip designs have been acknowledged for many years in mainframe design, and it accounts for IBM's controversial decision to set the clock speed of the 80286 chip in the original AT at 6 MHz, despite the fact that the same chip was capable of operating at higher speeds. IBM doubted the reliability of the supporting chips under the stresses imposed by higher speeds. When 33-MHz systems based on the 80386 first were introduced, many companies had problems getting them to work reliably. An associated problem with higher frequency chips is that, at higher speeds, the wires in the system might start to hum and resonate, creating "harmonics." This can affect the integrity of the circuits, resulting in unpredictable behavior. For mission-critical applications, you might want to err on the conservative side and avoid the newest, fastest chips until their reliability in the field has been proven.

The wave

Another important area in which hertz are used is the measurement of the alternating current that flows through the mains. This current goes from positive to negative and back to positive many times a second. In U.S. current, the complete cycle is carried out 60 times a second (60 Hz). In Britain, the mains current is 50 Hz. You might wonder why there is a difference in frequencies between the two countries. The answer has to do with what happens when electrical current is raised or lowered. The stronger the current that is being alternated is, the harder it is to control the change. You might think of putting a kink into a garden hose to turn the flow of water off and on. The more pressure in the hose, the harder it is to cut off the flow quickly and cleanly. Alternating 240 volts of current 60 times a second would make the cycles more abrupt than they are at 50 cycles.

Measurement of a smoothly alternating current should produce a smooth graph, like the first one seen in Figure 5.4. This is called a sine wave, because the cycle follows the mathematical curve generated by plotting the sine function. Abrupt cycles appear more like square waves and electricity following this pattern is more likely to create disturbances in electrical equipment. A phenomenon known as *harmonics* can occur at the "shoulders" of a square wave—frequencies related to but higher or lower than the normal frequency. Harmonics are a primary source of electromagnetic interference (EMI), which can seriously disrupt electronic components. An engineer using an oscilloscope can check the waveform of your current.

Fuses, Grounds, and Breakers

To understand some of the things that can go wrong with electrical circuits, it might help you to think of electricity coming from the power station to the outlets in your office, borne on the live wire, then returning to the power station on the neutral wire after it has performed its work. Materials along which electricity flows freely, such as the copper in your office wiring, are called conductors. Electricity essentially is lazy, attempting to get back to the power station as quickly as possible through any available conductor.




Figure 5.4 Current waves as sine, square, and modified square

Insulation

What keeps electricity from getting back too soon is insulation, which restricts the flow of electricity. Rubber, plastic, and a large variety of nonmetal materials make good insulators. Insulators make it possible to contain the flow of electricity, enabling us to make electrical appliances out of materials that are good conductors. For example, the casing of many personal computers is made of conductive metal, but touching the case should not give you a shock because insulators keep the current within the internal components of the system.

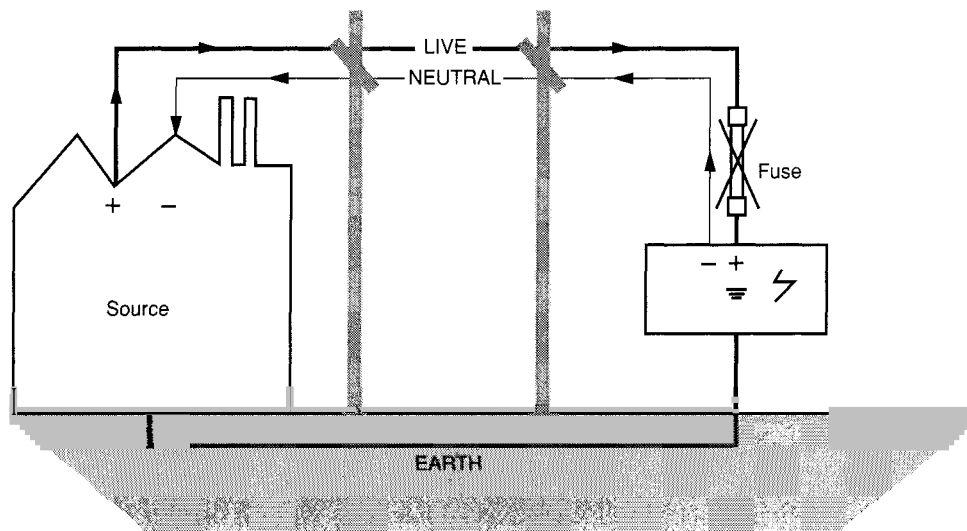
However, under extreme conditions, such as high voltage, even good insulators will break down, allowing electricity to flow where it should not. Faults in electrical circuits usually are caused by either an insulator or a conductor not working properly. Faults usually result in high temperatures, and the results can be disastrous for sensitive equipment. There are ways to prevent faults and there are technologies that minimize the impact of faults.

Grounding

Two alternative paths for the completion of an electrical circuit are down the neutral conductor of the supply cable or via the earth. In simple terms, the latter path exists because the neutral side of the electrical supply from the power station is connected to earth. This is why, if you touch a live conductor and are attached to the earth with no insulation, you will get a severe shock. The current passes through you to the earth.

The tendency of electrical current to return to earth has positive uses, called *grounding*. Grounding directs live current to the ground harmlessly, using the tendency of electricity to seek the fastest way home. Grounding is one way of dealing with electricity that has escaped from its insulation within an electrical device. To ground a device, all conductive materials in the device that should not become live, such as the metal case of a personal computer, should be connected to the ground. If insulation breaks down, current is diverted to the ground. In the process, the leak of electricity to ground should be detected, and the equipment turned off to prevent further damage. This task usually is accomplished by a fuse or circuit breaker, as diagrammed in Figure 5.5.

Electricity leaking from your computer onto the casing could be conducted to the ground in several ways. In some countries, such as Britain, where voltage is quite high, all appliances must be fitted with a three-pronged plug. There is one prong each for the live and neutral wires, and a third for the earth or ground. In the wiring of sockets, the wire from the ground prong is connected to a special cable called an earth continuity conductor. Several such conductors might be in a house or office. The metal casing of light sockets and any metal plumbing pipes also will have earth conductors attached to them. All such conductors will lead to a central point. Here the earth cable might literally go into the ground. Newer American buildings are designed in the same way, with the third prong leading to earth.



When electricity leaks to ground it causes a sudden increase in flow which blows the fuse




Figure 5.5 The process of grounding electrical equipment




Figure 5.6 Plugs and sockets in U.S. homes and offices

An alternative to running the cable into the ground, called protective multiple earthing, connects the central earth of the building to the neutral wire of the electricity company's service cable. This method is the basis for grounding when you only have two prongs on your plug. The neutral side can be used as the ground. However, this only works if the plug is put in the right way around, as diagrammed in Figure 5.6, and if the sockets are wired consistently.

This is why many American two-pronged plugs have one prong larger than the other and why receptacles will accept the plug only when it is correctly aligned (just so that we are clear on this, the small prong is the live side). Failure to follow the correct orientation can be dangerous.

Suppose you are in an older building and do not have a three-prong outlet handy for your computer. You might be tempted to place a "cheater" or adapter that allows you to use a two-prong outlet for a three-prong plug. This alone is bad enough as you are giving up the grounding of the computer. (Some adapters feature a wire that can be attached to a screw on the socket, which might establish a ground.) However, if you accidentally get the plug reversed, the results could be serious. If the computer has been designed to use the neutral wire for grounding, and the neutral wire now is the live wire, a leak of electricity with the computer could be disastrous. Grounding with personal computer equipment is part of the designer's job but, as you can see, it is possible to thwart the design. Fortunately, you also can take steps to help it work effectively.

Faults, fuses, and breakers

In the wiring of the computer, the casing normally is connected to the third prong of the supply cable. In some cases, the ground also might be connected to the neutral wire. If electricity should leak through insulation to the computer casing, then it will run straight through the earth conductor to the ground. At the same time, this leak of electricity will increase the current being drawn by the circuit. This increase should be sensed by a fuse or circuit breaker, two devices that are designed to interrupt a circuit if it becomes overloaded. (A fuse must be replaced after it has blown, whereas a circuit breaker can be reset after it has tripped.) You can see the typical deployment between such devices diagrammed in Figure 5.7.

When the live and neutral wires of a circuit are connected before the current has completed the work that the circuit was designed for, then you have a *short circuit*. This is an unexpectedly heavy flow of current that, like a flow of current to earth, generates heat. Because a short circuit or leak to ground indicates some form of component failure, you want the system to be turned off before any further damage is done. A fuse or circuit breaker, either in the computer or in the supply wiring, should accomplish this. Some computer peripherals, such as printers, have their own fuses. The multiple outlet extension cords widely used for personal computer systems often contain fuses or circuit breakers. In some countries, such as Britain, the appliance power cord plug has a fuse, as do most sockets, as diagrammed in Figure 5.8.




Figure 5.7 Deployment of fuses and circuit breakers.




Figure 5.8 British 240-volt mains system

How much damage is done before the fuse blows will depend upon the size of the fuse and its location. The further away the fuse from the appliance, the greater the potential damage to the appliance. Fuses are rated in amps; lower ratings, such as 1 amp, **will** blow under lower loads than higher-rated fuses, such as 13 amps. You should not replace a blown fuse with one of a higher amp rating unless you are absolutely sure that the blown fuse was rated too low.

If a piece of computer equipment blows a fuse or trips a circuit breaker, you first should turn off the equipment, then unplug the power supply cord leading to the equipment, and find the fault that caused the fuse to blow. Remedy the fault, then plug the equipment back in. Turn on the equipment, but be prepared to turn it back off again quickly if the fault was not properly remedied.

Among the more benign causes of blown fuses and tripped circuit breakers is overloading of a wiring circuit. For example, if you have an extension strip providing multiple sockets from a single wall socket, you could plug in equipment that draws more current than the single wall socket was designed to carry (see Figure 5.9). To correct this condition, you will need to rethink your plugs/socket arrangement, spreading the load more evenly.

More serious faults include damaged cables in which the insulation between conductors has broken down. Within appliances, insulation can wear down or melt down, leading to short circuits. Exercise great care when replacing fuses within computer equipment. All equipment should be turned off and disconnected before the fuse is removed. Bear in mind that some pieces of equipment, such as monitors, can hold a high voltage charge even after they have been turned off.

You must make sure that the replacement fuse is the same capacity as the blown fuse. For example, if the blown fuse is marked 2 amps, you should not replace it with

a 3-amp fuse. A 2-amp fuse means that the equipment will not be subjected to current above 2 amps. A 3-amp fuse will pass through 1 amp more current than the designer of the equipment intended. If a device keeps blowing a fuse, then something is wrong with the equipment.

Extension cords and capacities

Personal computer systems sometimes gobble up electrical outlets. Consider the following configuration, which is by no means untypical and which takes up six outlets: personal computer system unit, monitor, printer, scanner, transformer for mouse, and transformer for external modem. If you add this to the desk lamp, fan, and sundry other devices at the same location as the computer, you get a big demand for extension cords and outlet multipliers.

The use of such cords needs to be carefully monitored by office managers. Not only do the cords look unsightly, but they can present considerable danger to those who are apt to trip over them. Apart from the physical damage that can result from a sudden tug on the cord, it is a quick and nasty way to power down an entire system. For safety and security, consider implementing the following rules :

- Always match extension units to the load. Most have an amp rating on them, and the total amperage of appliances supplied by the cord should not exceed this figure. Extension cords should be kept out of pedestrian areas whenever possible. Use proper rubber mats to cover cords if they must cross pathways. Do not daisy chain extension sockets because this can result in more current than the cords are designed to carry. Use wall sockets whenever possible. Choose extension units that have fuses or circuit breakers in them. This will help limit any damage from electrical faults.
- Choose extension units and socket multipliers that have surge suppression built in.




Figure 5.9 Avoidable electrical connection errors

Regulating the Power Supply

Let us suppose that you have taken the basic security measures of plugging in your computer equipment correctly and making sure that it is grounded properly. Next, you need to make sure that the electricity flowing through the wires is free from interference, fluctuation, and complete disappearance. This is easier said than done. According to IBM research, in a typical month, a typical computer is likely to experience more than 120 power disturbances of some kind. The number of disruptive or destructive disturbances at the average computer site in the U.S. was put at 443 by the National Power Laboratory, a division of BEST Power Technology, a leading supplier of corrective devices. Some sources hold surges and power failures responsible for 45% of all preventable data loss (see Figure 5.10).

Terms and conditions

You can buy devices that will help you deal with these disturbances and prevent or minimize losses. However, you need to know something about how these devices work and what they do before you can decide which ones are appropriate to your circumstances. You also need to be familiar with the terms used to describe the problems that such devices address. Each of these will be discussed in further detail later in the chapter, but here are the short definitions.

Interference. *Interference* consists of spikes and noise. A *spike* is a sudden electrical impulse and an instantaneous and very dramatic increase in voltage. *Noise* is electromagnetic interference (EMI) and radio frequency interference (RFI).




Figure 5.10 Causes of preventable data loss according to consultants, Contingency Planning, as published in *PC Week* (1994).




Figure 5.11 Power disturbances according to Bell Labs

Fluctuation. Fluctuation consists of surges and sags. A surge or overvoltage is a short-term increase in voltage. This differs from a spike or transient in that it lasts longer, at least $\frac{1}{120}$ of a second, but involves a smaller voltage increase. A sag is perhaps better known as a brownout, a short-term decrease in voltage. According to a study by Bell Labs, sags account for 87% of all power disturbances (see Figure 5.11).

Disappearance. When the power goes completely, you have a power outage, otherwise known as a blackout or power failure. This can last for a few minutes or a few hours, depending on the circumstances. Common causes are excess demand on the grid, lightning strikes, construction work, high winds, and ice on the lines.

Sags

In the United States and Canada, mains current is nominally 120 volts. In the United Kingdom, it is nominally 240 volts. In France, it is nominally 220 volts. Most other countries are on one of these standards. The term nominal is significant because the actual voltage can vary. While U.S. voltage sometimes is quoted as being 110 V or 115 V, the voltage actually delivered to your office can easily be from 105 V to 125 V. The voltage in Britain usually is listed as 240 V, but actual measurements often range from 230 V to 250 V. Parts of Ireland use 220 V. Equipment rated at 220 V usually will work fine at 240 V (but might not survive an over-voltage condition of 250+ volts).

Fortunately, a lot of personal computer equipment can tolerate a range of voltages. A typical personal computer designed to run on 110 V to 120 V probably will work acceptably on anything from 85 V to 135 V. What causes problems for personal computer equipment are large swings in voltage (e.g., in America, a sag below 80 V and a surge above 140 V). What sort of problems? If a sag lasts for more than a frac-

tion of a second, the power supplied to the RAM (random access memory) will be reduced to the point where data is lost or perhaps scrambled.

Any work in RAM that has not already been saved to disk is at risk from a sag, but a sag also can affect data that is in the process of being written to disk. These days, to improve performance, many systems use a disk cache, which is a section of RAM that temporarily stores data being written to, or read from, a disk. This cache might be within system RAM or in RAM chips on the drive controller or the drive itself.

The effect of current returning to full strength after a sag also can have a disruptive effect. The long-term effect of sags is component failure, particularly in the case of drive motors. Symptoms of sags include unexpected system resets and "frozen" keyboards. Causes of sags usually are excess demand for power. This might occur locally, for example, when an air conditioner, elevator, or refrigerator starts up. Within some regions, the power company might systematically lower voltage at certain times of the day to meet demand.

Surges

The effects of a surge are harder to predict, and they depend to a certain extent on the power supply unit within the computer. The power supply has a moderating effect on the rise in current, but it might not be enough to prevent temporary shorting of components leading to scrambled data or even damage to chips. Erratic performance is the primary symptom of a surge, and premature component failure is the likely long-term effect. While sags can be caused by turning on electrical devices, surges often result from heavy equipment being turned off. A sudden excess of current is placed on the line.



Self Test I—Voltage

A typical digital multimeter, available from any electronic store, will give you a readout of voltage when you insert the test probes into a socket. In Figure 5.12, you can see a multimeter being used to check the current from a wall socket. An alternative is to disconnect the power supply cord from the back of the PC and insert the test probes in the openings of the power cord. This usually is more convenient than testing the actual wall socket, and most personal computer equipment has similar detachable power cords. Some meters take a moment to stabilize, but the readout should register a steady number after a few seconds. (Also note that many digital multimeters will tell you which wire is live and which is neutral, allowing you to check the integrity of the mains circuits that you are using.)

If the voltage readout continues to fluctuate, note the high and low measurements. If these are within six percent of the expected voltage, they probably are not going to cause a problem. If the swings are outside this range, it would be worth getting an electrician to check your wiring, then invest in some power conditioning and protection equipment.




Figure 5.12 Checking the voltage at the socket.

Spikes

A more dangerous and difficult to measure variation in current is the spike. This is a very sudden rush of current to very high levels. Mains circuits can carry damaging impulses as much as 2 kV (2000 V) that are very brief—a few microseconds in duration. Several different events can give rise to spikes. Many are caused by the switching on and off of large electrical appliances. Spikes come in two different varieties: normal mode and common mode. *Normal mode* events can be measured between the hot wire in a building's electrical circuits, and the neutral wire. *Common mode* events are measured from the neutral wire to ground. A normal mode spike of high magnitude can affect the power supply of the microcomputer. However, a common mode spike of only a few dozen volts also can blow out logic circuits or induce errors.

Surge/spike suppressors

A relatively inexpensive cure for surges and spikes is a surge/spike suppressor. (*Note:* The term *spike* is less widely used these days, and most products marketed as surge suppressors also handle spikes.) This is an electrical device placed between your personal computer equipment and the mains supply. Containing special circuitry that clamps down voltage when it starts to rise beyond an acceptable level, a




Figure 5.13 Surge and spike suppressors

suppressor prevents dangerously high mains voltage from reaching equipment. The circuitry in a suppressor is fairly compact, so these units can come in many shapes and sizes. Two examples are shown in Figure 5.13.

What do you look for when buying a surge suppressor? The first consideration has to be the reliability of your source. Do you trust the dealer that you are buying from? If the manufacturer is selling direct, can you get references from other users? Check the personal computer press for product reviews that include realistic bench tests. Beyond this, you need to look for suitable features and configuration. Do you want a unit that acts as an extension cord or do you want to insert the unit at the wall socket? You also should consider the following features.

The ratings. In the U.S., an organization called Underwriters Laboratory (UL) tests electrical equipment to make sure that it meets certain minimum standards, each of which has a number. The applicable rating for suppressors is UL 1449, and you should avoid products that do not meet this minimum standard. Another rating is the IEEE 587A, which measures "Let-through-Volts" when 6000 volts is fed to the suppressor. Obviously, you want the unit to let through as few volts as possible. A typical rating is 300, but it is possible to get suppressors that let through less than 40.

The ability to block spikes from reaching your equipment sometimes is measured in joules. A *joule* is a measure of energy: power expended over a period of time. For example, a product might be advertised as suppressing up to 140 joule spikes. You also will see measurements in amps, such as "140 joule spikes at 6500 amps." In general, the higher the voltage/joules/amps that the spike protector can block, the better.

You might see several other measurements in product literature for suppressors. A basic measurement is capacity, in terms of the total current that the device is designed to protect. This measurement works the same as for a simple extension socket. If the socket or suppressor is rated at 10 amps, then the total draw from all the equipment that you plug into it should not be more than that.

The term initial clamping *voltage* refers to the level of current at which the clamping effect of the suppressor's circuitry is applied. For example, in the U.S., this might be 140 V. Clamping response time is the length of time that it takes for the suppressor's clamping circuitry to take effect. This could be stated as nanoseconds, with 1 or 2 nanoseconds being good. The response time also might be stated in terms of cycles. Bear in mind that the ac mains in the U.S. alternates at 60 Hz, or 60 times a second. This means that 1 cycle lasts $\frac{1}{60}$ of a second, or 16.66 milliseconds. A response time of $\frac{1}{60}$ of a cycle would thus be 1.66 milliseconds, far longer than 2 nanoseconds. Response time sometimes can be distinguished from detection time, which is how quickly the suppressor detects that there is an overvoltage. Some makers argue that a short response time is of no use if detection time is too long. Obviously, makers will play up those measurements that reflect most favorably on their designs.

Breakers, failures, and insurance. There is some debate as to whether surge suppressors should be fitted with circuit breakers. As described earlier, a circuit breaker is a resettable switch that turns the supply off if the circuits become overloaded. For some people, this is the minimum level of protection and is found on multiple outlet extension boxes that don't have specialized surge-handling circuits. Note that fitting a circuit breaker to an extension box does not make it a surge suppressor. You need additional circuitry to effectively protect against surges. Also note that you should not reset a circuit breaker that has been tripped until you have determined what caused it to trip in the first place.

Some surge suppressors do not have circuit breakers but are fitted with monitoring lights that indicate surge protection is active. If the light goes out, the surge system has failed and the unit needs to be replaced. However, on some cheaper units, the fact that surge suppression is no longer in effect does not disable the unit, meaning that surges will be passed through to equipment. When an inexpensive grocery or drug store surge suppressor, such as the GE SurgePro (about \$7 for a three-outlet unit), needs replacing, this means you buy another one.

Surge suppressors from some specialist manufacturers, such as APC (American Power Conversion), come with a lifetime replacement warranty, so the failure of the unit is not going to cost you. More expensive units like this (\$25 or more for a three-outlet unit) shut down when they are no longer capable of providing protection, which is a safer approach. Some manufacturers back their systems with insurance, so that you might be able to get reimbursed if equipment is damaged while their suppressor is being used.

Additional protection. Many surge suppressors provide multiple sockets for plugging in your equipment. The design of the unit should be such that each socket is protected separately. This design is more likely to successfully clamp heavier surges than a design that simply protects the single line running to multiple sockets. Sepa-

rate protection also can contribute to the reduction of noise interference between pieces of equipment plugged into the same supply circuit (see the following section for more on the noise problem). If you are going to spend money to protect against surges, you probably will want to get a unit that also has spike and noise protection. Protection against sags typically is provided by supplementing the mains voltage with battery power. This task is performed by an uninterruptible power system, or UPS, which will be discussed in a moment.

Because electricity is not fussy about which wires it travels down, it is possible for surges, spikes, and noise to affect phone lines, cable television cables, and network wiring. It is likely that one or more of your computers is attached to one or more of these circuits. You can buy surge suppressors specifically for these circuits or purchase a mains surge suppressor that also handles other wiring (see Figure 5.14).

To give you an example of how this works, the computer that I am writing on at the moment is plugged into a surge suppressor. This is plugged into the mains. My internal modem is plugged into a special port on the surge suppressor. The phone line from the wall socket is plugged into another port on the surge suppressor. My fax machine, which shares a line with my modem, also is plugged into the surge suppressor. In my living room, the power cords for my TV and VCR are plugged into a surge suppressor, along with the coaxial cable from my cable company.

An alternative approach, which might be more convenient in a larger office setting, is to use separate devices for each type of circuit. For example, APC makes surge protectors for phone lines (with RJ-45 sockets), RS232 serial ports (DB-9 and DB-25), and networks (Ethernet and Token Ring). You only have to lose one computer or a few hours worth of data to a surge on an unsuppressed circuit to know that the protection is well worth the money.

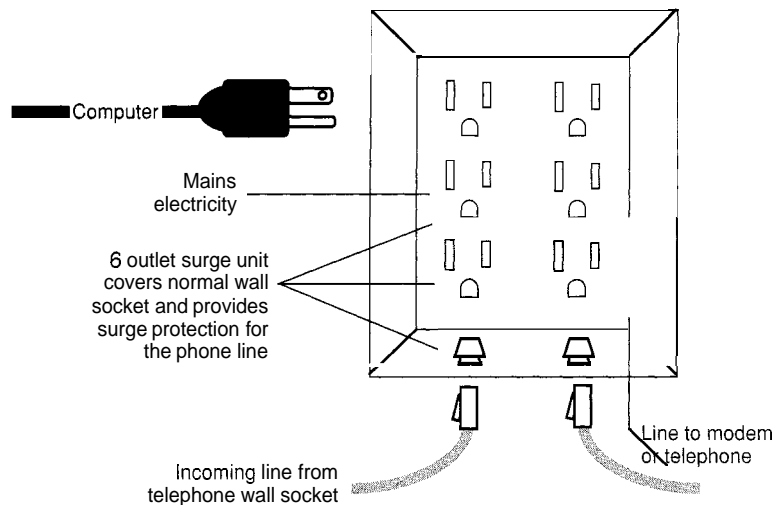


Figure 5.14 Surge protection goes beyond mains circuits



Lightning Strike

The damage that was done was 400 yards away from the actual strike. The lightning traveled from a phone socket at a house that was hit, down the underground phone wire, and out through a wall socket at the house of the "victim." From there, it went into his modem, out along the serial cable, and into the PC via the serial port. After destroying the motherboard, video card, and parallel port, it traveled out through the parallel cable into a print sharer. At this point, the lightning split two ways, taking out another computer and a laser printer. Back at the original PC, it also traveled out through a teletext card and blew up a TV antenna booster. A fax machine plugged into a second phone line also was destroyed.

Report in *Personal Computer World*, 1994

Switching problems

Earlier you saw that switching electricity on and off presents difficulty for chip designers as it is difficult to raise voltage from zero to a set level, then reduce it back to zero in a very short period of time. Switching current is not only a problem for computer designers. Anytime an electrical circuit is completed or broken, it can have negative side effects. To understand this, think of electrical flow as a water supply that flows through a pipe until it meets a tap. The effect of a tap is like that of a switch in that it turns supply on and off, but a tap acts differently from a switch in that it operates gradually. Turning on the tap gradually opens a valve that slowly increases the flow of water. Electricity also flows under pressure, but most electrical switches do not act like valves. Switches act more like jamming the flow of tap water with your thumb to turn it off, which usually results in water escaping around it, spraying out under pressure as you try to seal off the flow.

When switches are turned on or off, some of the electricity still might "spurt out" in a spark, short, surge, or spike. If you turn a hair dryer on and off in a dark room, you probably will see this phenomenon. If you unplug the hair dryer when it still is turned on, you probably will see a spark at the socket.

These flashes of current can have two negative impacts for sensitive personal computer equipment. First is the spike, the sudden rush of voltage that spike protectors defend against. You can help cut down on this problem by following a few rules. Do not plug or unplug appliances that are turned on, particularly computers, printers, and monitors. Often these appliances will have some form of protection on their switch circuits that you bypass if you leave the switch on, then plug it in or unplug it. Because separately switching on each piece of equipment is a chore, you might want to use a power center, a surge/spike protection unit in a console design that pulls together the supply for each part of the system and can handle the current flows needed at system start-up. You can see an example diagrammed in Figure 5.15.

The second negative effect of switching is the more complex matter of harmonics—electrical frequencies substantially higher than the current that produced them. The sudden action of a switch is like the sharp flick of the finger that produces harmonics from a guitar string. Emanation of these unintended frequencies




Figure 5.15 Diagram of power-switching unit

from one piece of equipment can interfere with the operation of neighboring equipment. Good surgespike protectors that supply electricity to more than one piece of equipment will offer some form of isolation for each piece to cut down on this problem, which sometimes is described as "noise" and will be discussed in more detail in a moment.

Sags and Line Conditioning

Surge suppressors do not prevent sags, and they do not protect against prolonged overvoltage conditions. If you are operating computers in an environment that is particularly harsh from a power supply perspective, you might want to invest in a *line conditioner*. This is a device placed between equipment and mains supply, like a surge suppressor, but it does more than clean up surges, spikes, and noise. A line conditioner puts out constant power within tight specifications despite chronic line voltage variation. A good line conditioner actually can boost undervoltage to counteract minor sags.

Beware of products that are labeled line conditioners but are no more than surgespike suppressors. While a good line conditioner offers even better protection against surges and spikes than an ordinary suppressor, it should do more than that. A line conditioner should be capable of monitoring the incoming voltage and adjusting it either up or down so that the output voltage falls within a narrow band (such as 103 V to 132 V in the U.S., as specified in ANSI C84.1-B). If you are worried about frequent sags and complete power outages, you should be thinking of getting some form of backup power device. These are described in a moment and they also offer surge and spike suppression.

Noise and Static

Surges, sags, and spikes are not the only electrical supply problems facing personal computer users. There also is the matter of noise—not the sort of noise that you can hear, but the electronic noise that interferes with the performance of electronic components. To understand this problem and place it in perspective **will** take a few paragraphs, but understanding noise will help you to stop it from interfering with your data and also will prepare you to combat electronic eavesdropping on the noise that your personal computer emits (a problem described later in this chapter).

Terms and concepts

Two terms are used to describe this noise: *radio frequency interference* (RFI) and *electromagnetic interference* (EMI). You literally can see this noise if you use an



Self Test II—Hear the Noise

If you want to experience this for yourself, try using an FM radio to listen to your personal computer as it accesses the disk drive. First, set up a macro or batch file on your personal computer that carries out a continuous loop of disk access activity. On a PC, you might write a batch file like EMI.BAT listed here:

```
:START
ECHO Testing EMI
DIR/W > EMI.DAT
DEL EMI.DAT
GOTO START
```

This repeatedly reads the directory into a file called EMI.DAT, then erases the file, in a loop that will continue until you press Ctrl-Break or Ctrl-C. Now, tune the radio so that it is between stations, at around 93 MHz on the FM dial. You should get a fuzzy static sound (try using headphones—the noise can be quite annoying to others). While the batch file executes, slowly adjust the tuning on the radio until you hear a rhythmic buzz or crackle that matches the rhythm of the batch file. (Try tuning up from 92 MHz to around 96 MHz; the test works best in areas with relatively few radio stations, but you should detect a signal of some sort, even if it is interference with an actual station.)

When you do pick up the noise of the disk access, try moving the radio away from the computer to see how far away you can detect the signal (you might need to twist the radio through different angles to keep the signal). You probably will find that you still can get the signal several yards away. What you are hearing is EMI—radiated harmonics of the circuits accessing the disk. You can do a similar test for emissions from your keyboard. If you try typing while tuning your radio, you probably will be able to pick up keystrokes as blips or abrupt changes in the background static. I found that my Sony Walkman was able to detect keystrokes from a cheap PC keyboard at 93.2 MHz on the FM dial. The original IBM PC keyboard could be picked up on 95.4 MHz, with sufficient clarity to distinguish different keystrokes. With specialized equipment, it is possible to record such signals, analyze them, and determine what is being typed.

electric drill close to your television set. The electric motor in the drill causes lines, snow, or other patterns to appear on the screen. Similar EMI can be caused by unprotected spark plugs in a car outside. Radio interference can result from cordless telephones that use radio waves to communicate between the hand unit and the base. Not only TV reception, but also the integrity of data within personal computers, is at risk from these and other sources of interference. The mains supplying electrical current to your personal computer can act as a noise transmitter, carrying EMI and RFI into the computer system.

Whenever an electric current flows through a conductor, it generates a force field—waves of energy that create an electromagnetic field. The frequency of the waves in this force field can extend into the frequencies used by radio waves used to transmit radio and TV signal waves. You do not need high voltages to create EMI. The strength of an electromagnetic field is more directly related to the acceleration of the electric charges, the rate of change in the current, rather than the strength of the current itself.

Any conductor that carries an electric current acts as a broadcasting antenna radiating a field produced by that current. A typical electrical appliance can contain hundreds of conductors, varying from a fraction of an inch to several inches in length. Probably, at least one of these conductors will resonate with one of the frequencies generated by the electronic components within the appliance, turning the conductor into an antenna emitting signals that interfere with other appliances.

Noise protection

Good design can cut down on these electronic emissions or radiation. Metallic casings can be used as a shield to stop emissions from escaping. Computers that have plastic outer casings often have a thin metal lining inside. However, achieving complete protection is impractical. Furthermore, design flaws and faulty components can result in extensive emissions. You might have seen a notice like the one in Figure 5.16 in the manual accompanying such components as monitors and system units. This notice is required by the Federal Communications Commission whenever a product is sold that can interfere with a radio and television reception. To protect your personal computers from EMI and RFI, you need to consider three aspects: the power line input, placement of noncomputer appliances, and shielding of computer equipment.

Power line noise. Some surge and spike suppressors are designed with circuitry that filters out noise from the power supply. Noise suppression is measured in decibels, as in "offers high frequency noise suppression of 20 dB @ 50 kHz." While it is difficult to put the meaning of such measurement in practical terms, it is useful as a relative measure of effectiveness, a higher dB rating being better.

Placement of appliances. As a general rule, personal computers and heavy electrical equipment do not mix well. When setting up personal computer workstations, you should try to keep them away from such equipment. It is difficult to suppress interference from the powerful currents running through such machinery as overhead

RADIO FREQUENCY INTERFERENCE STATEMENT

This equipment generates and uses radio frequency energy and if not installed and used properly, that is, in strict accordance with the manufacturer's instructions, may cause interference to radio and television reception. If this equipment does cause interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

Reorient the receiving antenna.

Relocate the computer with respect to the receiver.

Move the computer away from the receiver.

Plug the computer into a different outlet so that computer and receiver are on different branch circuits.

If necessary, the user should consult the system owner's manual, the dealer, or an experienced radio/television technician for additional suggestions.

The user may find the following booklet prepared by the Federal Communications Commission helpful: "How to Identify and Resolve Radio-TV Interference Problems." This booklet is available from the following sources for \$5.00 postage-paid, price and availability subject to change:

U.S. Government Bookstore
World Savings Building
720 North Main Street
Pueblo, Colorado 81003
(303) 544-3142

Consumer Information Center-V
P.O. Box 100
Pueblo, Colorado 81002
(303) 948-3334

Figure 5.16 FCC Notice

cranes, printing presses, and electric welding tools. For most offices, this is not a problem. Other office equipment, such as photocopiers, usually are well shielded. However, elevators can be a problem in office buildings, and the industrial use of personal computers is growing rapidly. Therefore, in some situations, it might be necessary to place a metal enclosure around a personal computer to protect it from electrical noise interference. See the section on electronic eavesdropping for more on shielding.

Other computer equipment. A good surge/noise suppressor will filter out the noise interference between the various components plugged into it. However, the outer casing on some components might not be properly shielded, resulting in interference between devices. It is useful to bear in mind that noise incompatibility problems do arise from time to time, even among components from the same manufacturer. You might be familiar with the rule of external disk drive placement with older Macintoshes: always below the Mac or on the right of it. Interference from the Mac's power supply, which was on the left of the original Mac body, was so strong that it disrupted the read/write activity of the drive heads. If you have erratically performing equipment, you might consider rearranging it with respect to other parts of the system. This might just solve the problem.

3

Self Test III—Dr. Gauss

You can test emissions from your computer equipment with a Gauss meter. See appendix D for more details as shown in Figure 5-17.




Figure 5.17 The Dr. Gauss EMF Detector

Phones and faxes

Electrical noise also affects telephone transmissions. As mentioned earlier, you can get surge suppressors with built-in noise filters for phone lines that carry data and fax transmissions. In some cases, these are combined with regular surge/spike suppression.

The other side of the phone noise problem is interference from phones affecting personal computers. This often happened with the early cordless phones, and you still might need to keep a cordless unit away from the computer. Do not be surprised if speaking into a cordless phone while looking closely at your personal computer monitor causes some disruption to either the phone or the computer.

Coping with static

In some parts of the world, at some times of the year, static electricity is a major problem for personal computer equipment and data. This is not the place to go into the

physics of static electricity. Most people are aware that an electrical charge can build up in the human body, which acts rather like a capacitor. When the skin, a conductor, contacts another good conductor like a metal desk, the static charge can be released, sometimes at quite high voltages. If a strong static charge is grounded through your computer equipment or a floppy disk, then severe disruption of processes and data can result.

The antidotes to static are numerous. Dry air fosters static buildup, so maintaining proper humidity in an office will cut down on static (as well as improve sinuses). Badly affected offices can be fitted with special carpeting, rugs, and mats, which cut down static buildup. An antistatic mat can be placed under your chair to help prevent a charge from building up. Static grounding now has been designed into a variety of accessories. Small antistatic mats can fit under the personal computer keyboard. Touching the mat before typing or handling disks will ground the static in your body, leaving you safe. Monitor screens and mouse pads with similar protection also are available from many office supply houses.

Before investing a lot of money in antistatic products, bear in mind that static is not a universal problem. Unless you have seen or felt sparks around the office, you probably have nothing to fear. A rough test can be done by taking off a sweater in the office at night with the lights out (I would never suggest that this particular test be used as an excuse for inappropriate behavior). Walk around the office for a few minutes in stocking feet, then take off the sweater while keeping your eyes open. If you see sparks, then you know that there is static around, and some basic precautions are definitely in order.

Guaranteeing the Power Supply

There are ways of making sure that your computer does not lose power when there is a blackout. These vary considerably in price and capability. You will need to weigh these factors against the potential loss of data and damage, as well as inconvenience that a power outage can cause.

A measure of protection

The sags, surges, and spikes discussed earlier can have a negative impact on all types of electronic equipment, including personal computer system units, monitors, printers, and other peripherals. However, the negative effects of a complete loss of power primarily concern the system unit. Unexpected powering down of the system unit can:

Wipe out data in RAM. Newly entered data or recently edited data that has not been saved to disk is lost.

Trash the cache. Data that has been "saved" but is waiting to be written to disk is lost.

Interrupt disk writes. Important information required by the operating system, such as file location, can be lost, resulting in files being lost or scrambled.

Scramble encryption. If you are password-protecting data, a loss of power during encryption might render it irretrievable.

Crash the hard disk. The read/write heads of most hard disk drives automatically retract from the disk when the unit is powered down; however, in older systems, the heads can "crash" onto the disk surface and damage it, causing a loss of data and even physical damage to the disk.

Interrupt printing. When power is returned, uncompleted print jobs must be resumed. This can be a very tiresome chore unless you have good print management software. In some cases, the entire print job must be redone, after you have unclogged any paper jams caused by the outage.

Interrupt communications. When power is returned, data that was being transferred between computers must be checked for accuracy and files that were in the process of being transferred might need to be transmitted again.

Halt operations. In organizations that depend heavily upon computers, a power outage and the temporary lack of computer facilities could adversely affect productivity and profitability.

Cause spikes and surges. These can occur when power is returned. Normally, you will want to switch off computer equipment when the power goes out, but this is not always possible. When the utility company restores power, it often returns unevenly, which could damage appliances that were not switched off.

Clearly such events are to be avoided if possible. While power from the mains is susceptible to all manner of fluctuations, down to complete absence of current, batteries provide much smoother power, independent of the mains supply. However, batteries do not last forever and eventually will run down, at which point they are disposed of or recharged, depending upon their design. Given these factors, it would seem logical to power a personal computer from batteries that are recharged from the mains. This is a good idea in principle. You can buy equipment, known as an *uninterruptible power system* (or UPS), that will provide power in this way, offering protection from power outages.

When the mains power disappears, the battery in a UPS continues to supply power. This power is limited by the battery capacity in both duration and strength. Powering all of your computer equipment from a UPS during an extended power outage is not feasible (for this you need your own generator). Typically, the role of the UPS is to provide enough power to see you through brief outages of a few minutes duration and to enable you to carry out an orderly shutdown of your system during prolonged outages. The UPS system will warn you when the battery is getting low so that you can save files and power down. Some systems come with software that will manage at least part of the shutdown process.

The need for a UPS

There is no doubt that a UPS will improve the security of your data. It takes only one outage in the middle of a crucial project to convince you of the desirability of a backup power system. At a minimum, all network servers should be powered through a UPS. Novell also recommends that all workstations be UPS powered. Your decision to purchase a UPS might be influenced by the number of times you have experienced a power failure while computing.

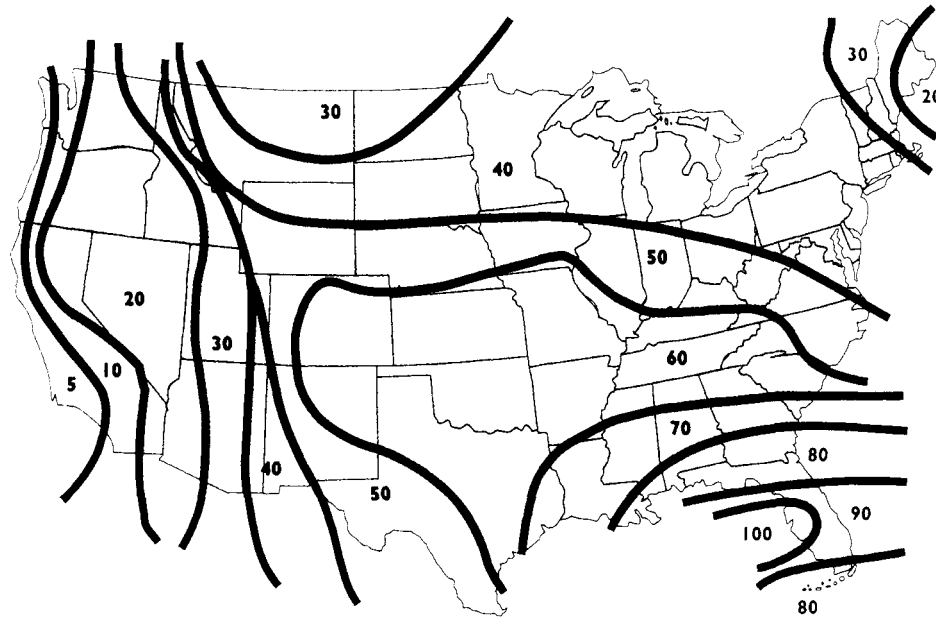


Figure 5.18 Map of the average number of annual thunderstorms in the US based on National Weather Service data

Obviously, some places get more unscheduled outages than others, with remote and rural locations generally faring worse. However, urban areas with heavy construction tend to experience an above average number of interruptions, as do areas, both rural and urban, where lightning is frequent (see Figure 5.18). Furthermore, power outages might well be on the increase, due to the predicted under capacity of many utility companies. In the U.S., spending on utilities has dropped from 2.3% of GNP (Gross National Product) in the 1960s to under 1% today. Areas such as the Southeast and New England might find themselves as much as 13% undercapacity in the second half of the 1990s.

A power outage scenario

Thinking about a typical power outage scenario in an office without a UPS might help you decide whether or not you need one. Some power outages are announced so that precautions can be taken. Most strike without warning. Typically, the loss of power in an office without standby systems is accompanied by a general groan and the odd expletive. Users who have not recently stored their data on disk have just seen their work disappear from RAM as well as from the screen, wasting valuable time that could run into hours. As the office manager frantically reminds users to turn off their systems to prevent damage from the surge of returning power, people begin to consider what they have lost.

Communications sessions, such as downloading data from other computers, have been broken. Print tasks have been aborted in midline. File transfers across net-

works have been disrupted. File save operations have been interrupted. There is the possibility of damage to hard disks. When power is returned, users will need to restart their systems and make sure that they were not damaged. They will need to check for data that might have been scrambled. Print jobs will need to be sorted out, completed, or resubmitted. Work that was lost will need to be repeated. The major impact of productivity thus becomes the recovery from the effects of the outage, rather than merely the computing time lost due to lack of power.

What is computing like when you have a UPS? A typical UPS will sit quietly in your office largely unnoticed while the mains power is on. A power cord leads from the mains to the UPS, and a variety of PC system units and monitors are plugged into the back of the UPS. Most UPS systems have one or more warning lights to let you know that the mains are okay and that the battery is charged or charging.

Suppose your mains power is cut off. An alarm in the UPS is sounded. A warning light shows that you are running on the battery, but the UPS continues to power the equipment. Following procedures laid down in the organization's security or operations manual, users react to the alarm by starting to save their work. This is not always a case of picking Save from a program menu. The computer might be in the middle of a lengthy operation, such as database sorting, and the user will have to decide whether the process can be completed within the time remaining before the batteries are depleted.

As work is being saved, the UPS reports on battery status, sounding a new alarm when only a few minutes of power remain. Suppose that the mains power is restored at this point. Work can carry on with virtually no interruption or loss of productivity. Any computers that already have been turned off can simply be restarted. What if power does not come back for an hour or more? The UPS alone cannot keep your operation going this long; however, it has enabled an orderly shutdown and minimized disruption and data loss. It is not difficult to see that, in either case, a UPS makes life a lot easier.

A word of caution

As you can see from the diagram in Figure 5.19, a UPS looks like a box with wires and lights. Indeed, while there are several different types of UPS, and hundreds of different products that are advertised as UPSs, they all tend to look like boxes with wires and lights.

The simple appearance of a UPS, combined with the fact that its performance is difficult to test without special equipment, means there is considerable potential for ripoff. In Figure 5.20, you can see just how complex a respectable UPS looks when you open it up. To avoid buying a UPS that does not live up to its name:

Be sure to read a technical review of its abilities first (you can find these in *PC Magazine*, *BYTE*, or other trade journals).

Look for a manufacturer that has ISO 9001 quality certification for its manufacturing and operations.

As soon as you have installed a UPS, test it yourself by turning off the mains yourself. Be sure to test the full load that will be placed on the UPS.




Figure 5.19 Diagram of UPS deployment

How a UPS works


Providing electricity to your personal computer at full mains strength is a difficult task for a battery to perform. Batteries produce direct current, while mains-powered PCs consume alternating current. Furthermore, batteries work more efficiently at lower voltages, such as the 12-volt systems in cars. Alternative power supply systems, such as wind-powered electrical generators, typically produce 12-volt direct current that is stored in batteries, then converted to alternating current when required by normal household appliances. The conversion from dc to ac is carried out by an inverter.

One approach to providing uninterrupted mains level power to a computer involves converting mains power down to 12 volts dc and using this current to charge a battery. The battery output then is inverted back to mains level when the computer draws from the battery. The entire chain of activity can be seen in Figure 5.21. This system sometimes is referred to by its advocates as a "true UPS" because, as far as the equipment connected to it is concerned, there is no change in current when the mains power fails.

If you consider what is happening in Figure 5.21, you might find it somewhat ironic. Most components of a personal computer system run on just a few volts of dc power (usually from 3 V to 12 V). The power supply unit with a desktop computer acts as a transformer to create low dc voltage from much higher ac current. Using batteries in a UPS to create 110 current that then is transformed down to 12 volts or lower might seem like a lot of trouble. Notebook computers demonstrate a much simpler arrangement, as shown in Figure 5.22. The system runs off battery power but is plugged into the mains most of the time. That way, the batteries are constantly recharged. The user can continue working for several hours after the mains power has failed. Furthermore, most notebooks have power monitoring software that warns the user when battery power is getting low. If you work in areas where power outages are a constant threat, a notebook computer might be a much better arrangement than a desktop backed by a UPS.

U N I T Y / I Single-Phase UPS

Microprocessor Brains Replace Brute Force



- ① UL 1449-Listed for surge suppression; line filtering for clean power
- ② Proprietary fault-tolerant inverter design for fail-safe operation
- ③ Microprocessor controlled, multi-mode, time-equalized charger for optimized battery life
- ④ Advanced self-diagnostics for continuous operation
- ⑤ Digital display with status LEDs and 5-button keypad for UPS control
- ⑥ Multi-tasking, high-speed microprocessor control outsmarts power problems
- ⑦ Phoenix[™] allows advanced, remote diagnostics
- ⑧ Automatic bypass keeps loads running
- ⑨ High redundancy tap-switching for superior voltage regulation and maximum uptime
- ⑩ Hot-swappable batteries eliminate maintenance downtime
- ⑪ Reliable multi-tap isolation transformer adjusts to incoming power
- ⑫ Isolated RS232 port for remote operation and communications

Figure 5.20 Inner view of a "real" UPS.




Figure 5.21 Diagram of power supply using a UPS.




Figure 5.22 A notebook computer running on constantly charged batteries

UPS terminology

Note that the acronym UPS stands for uninterruptible power system and not power supply. A UPS is an external system for providing continuous power when the mains supply fails (see Figure 5.23). A power supply is a component inside the personal computer that converts the utility's ac power to dc, which is what the logic circuits of the microcomputer need for working energy.

The correct terminology is even more complicated because not all UPSs are truly uninterruptible. Some experts divide UPSs into online or offline. While an online UPS is always drawing power from a battery, meaning that the backup supply is always "online," the offline system kicks in backup power only when the mains current fails. Some people would argue that an offline unit is more correctly called a standby power system (SPS). However, some SPS designs are virtually uninterruptible, further blurring the distinction. Indeed, APC uses the term UPS to describe all of its standby systems (see Figure 5.24).

One reason for this blurring of terms relates to the personal computers themselves. A personal computer power supply has what is called ride-through, which is the amount of time that the power supply can deliver stored energy to the logic circuits with no electricity being fed to the supply. This energy storage is directly related to the size and quality of the power supply components, particularly the filter capacitors. The ride-through of a typical personal computer is from 20 to 40 milliseconds. This is a long time in the terms of electronics because a single cycle of ac current at 60 hertz takes only 16.66 milliseconds ($1000 \div 60$). To put it in nontechnical terms, the electricity reaching the computer power supply probably could skip a beat without affecting the flow of power to the computer's logic circuits.

The SPS solution

This long ride-through has enabled the SPS to become the most popular of the power loss protection devices for the personal computer market. In Figure 5.25, you can see a block diagram of an SPS. The incoming mains power (ac) is fed directly into the microcomputer under normal conditions. When mains power fails, the transfer switch senses this happening and turns on the inverter, which converts battery power (dc) into an ac source that keeps the computer running. When mains power returns, the transfer switch returns the computer to mains power. This is called standby technology because the inverter is literally "standing by," waiting to be turned on.

The UPS solution

In Figure 5.26, you can see a block diagram of a true online UPS in which the incoming utility power is converted from ac to dc by a rectifier/charger. As the name implies, this component performs two functions. The functions are changing the power to dc (the rectifier) and charging the battery. The battery is said to float on a dc bus, meaning that a single conductor connects the rectifier/charger and the inverter. If the battery needs charging, it draws power from the bus. If, on the other hand, the bus voltage level falls below the battery float voltage, the battery delivers energy to the bus. The energy conducted through the dc bus provides power to the inverter, which then provides power to the microcomputer. In other words, the system is online all of the time, and a full-time ac-dc-ac conversion takes place.




Figure 5.23 A range of UPS models from BEST Power Technology




Figure 5.24 A range of UPS models from APC

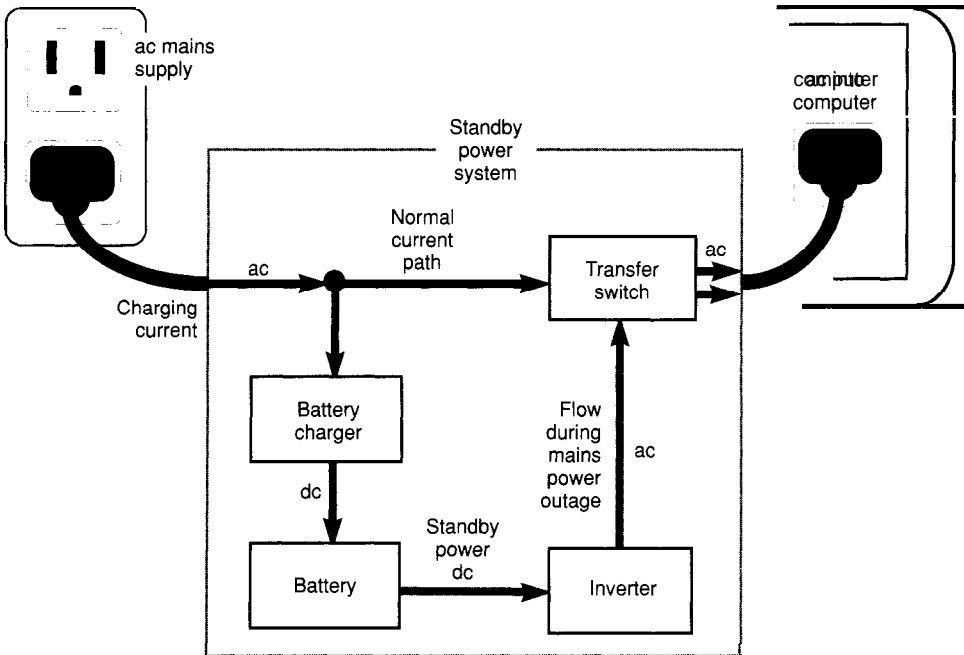


Figure 5.25 Diagram of a standby power system.

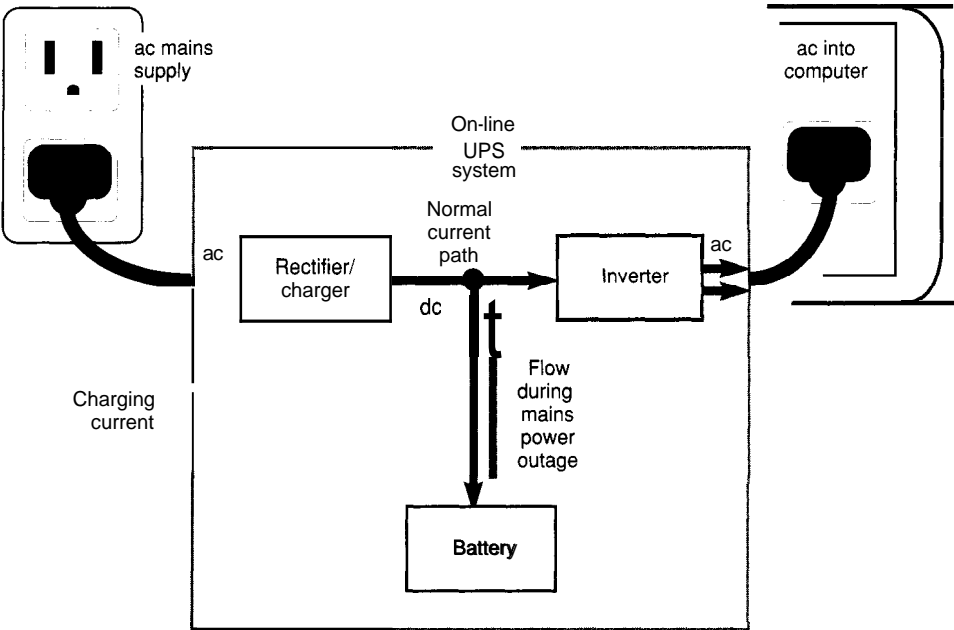


Figure 5.26 Diagram of an on-line UPS.

The advantage of this design over a standby system is that no switching takes place if utility power fails. Because the inverter is always providing power to the personal computer, it never sees an interruption of power. This feature is not without its drawbacks. The duty cycle of the components, the percentage of time that they have to work, is 100%. This means that they must be bigger and with higher reliability ratings than the equivalent components in an SPS of similar output. Particularly affected is the battery, which is working all of the time in a UPS. The battery in a UPS is likely to wear out sooner than that in an SPS. An online design can cost considerably more than a standby unit of the same rating.

A special transformer

One technology, the ferroresonant transformer, offers a middle ground between the expensive but ever-present power of the UPS and the cheaper standby systems. Interestingly enough, ferroresonant transformers have been around for a long time. They have the unique ability to store energy for a few tens of milliseconds. In Figure 5.27, you can see how this device can enhance the performance of the simple SPS.

Notice that all of the blocks are the same but that the transformer has been added at the output. With this design, the time that it takes to switch on the inverter is covered with the ride-through capability of the transformer. Because the personal computer's own power supply unit also has a ride-through factor, the computer is unaware that any switching has taken place. Companies, such as BEST technology, pushed this type of system as a good price/performance compromise between a true UPS and a regular SPS. As you might expect, a ferroresonant transformer does tend to add cost to the basic SPS, but this is partly offset by the fact that it adds some real power conditioning to the current that it supplies. You can expect to see an increas-

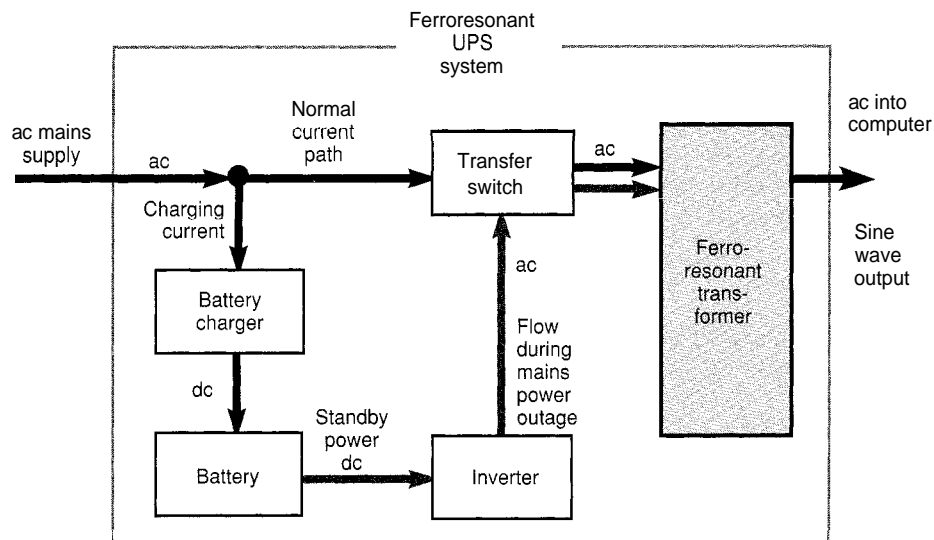


Figure 5.27 Diagram of an enhanced SPS/UPS

ing number of SPS units with ferroresonant technology, priced between the UPS and the regular SPS units.

Choosing your SPS/UPS

The variations of the uninterruptible designs on the market are virtually limitless. Expect to see hybrids of the technology discussed earlier, plus innovations and exceptions. However, a working knowledge of the basic designs shown here will prepare you for choosing a design for your own purposes. There is considerable argument between proponents of the various UPS technologies. Supporters and manufacturers of online models claim they are better because the inverter is online all of the time. They point to thermal stress that can affect inverters when they are started from cold. This is not a worry for an online system because the inverter is always on. Fans of standby designs say that they are more reliable because the inverter is only on when power is out and thus subject to less wear.

To the consumer seeking reliable protection for data, the main concern is not whether to side with one point of view or another. The question is how well the system is designed and constructed. If an SPS is not designed to pick up the load gracefully, even the most reliable components **will** not cope. If an online unit is not carefully engineered, then thermal stress **will** cause premature failure while the unit is running. The inverter is not necessarily the weak link in a UPS. More important factors are quality control and batteries. Any unit that you buy should have a test certificate and a long warranty, including free replacement if the system fails within the first year of service.

Batteries are the single most frequent point of failure for any UPS, so you should check the battery source and specifications thoroughly. Models with "hot-swappable" batteries should be considered (this means that the battery can be safely changed without interrupting the power supply). With some idea of how they work and what level of protection they provide, it is time to assess whether you need a UPS or an SPS, and if you do need one, how you select one that is suitable for your needs. You should look for several important features when purchasing a UPS.

Warning lights and PC interface. You need some way of knowing what your UPS is doing. Many designs feature indicator lights that let you know when the battery is being used and what kind of charge the battery is holding. Look for indicators that are functional rather than just decorative. You probably will want audible as well as visible warnings when the system switches to batteries and when batteries are low.

Many systems now offer a personal computer interface—a means of sending signals to a personal computer to warn of impending shutdown. This is particularly useful if you are using a UPS to protect a network file server. For the feature to be of any use, the personal computer must be running software that can interpret the signals sent from the UPS. Various programs can do this, and they are described in the context of network security in chapter 12. A good example of this type of software is PowerChute from APC, which allows you to test and monitor the UPS, execute a safe system shutdown in the case of a severe outage, and run diagnostics remotely (this requires a UPS with a suitable communications connection).

A switch in time. A major factor in the performance of any SPS or offline UPS is switching time. The unit must be able to switch its inverter on and smoothly—some would say gracefully—take up the electrical load before the personal computer's internal ride-through expires. Typical switching times are from 4 to 10 milliseconds. Times within this range mean that the computer's power supply will not notice any change in current. However, when checking product specifications, you should make sure that quoted switching time includes the time that it takes for the SPS to sense an outage and complete the switching process.

Adjustable transfer point. An SPS or offline UPS will switch to battery power according to a predetermined transfer point or voltage level. When the mains voltage (say 120 V) falls below the transfer point (say 100 V), the SPS begins the switching process. This ensures that, by the time voltage reaches a dangerously low level (say 80 V), the microcomputer already will be on battery. A selectable transfer point is an important feature. The power supply inside most microcomputers has a working voltage window that is enormous, from about 80 V to nearly 140 V. If your site experiences chronic brownouts or low-voltage conditions, you might want to buy a unit that lets you select a low transfer point (possibly as low as 90 V) so that the SPS does not transfer to battery power unnecessarily.

Hysteresis. The voltage at which the SPS switches back from battery power to mains when the latter returns is called the *retransfer point*. You should be able to set this above the transfer voltage so that, if the utility voltage hovers near the transfer point, repeated switching on and off of the battery does not occur. This feature is referred to as *hysteresis* (literally meaning that the effect lags behind the action that causes the effect). A typical hysteresis window might have a low of 102 V and a high of 107 V.

Waves. The ac current from the mains has a smooth sine waveform. Many SPSs do not put out a sine wave because it is considerably cheaper to put out a square wave, a rectangular wave, or some quadrilateral in between. The inverter that creates ac from dc is basically a very fast switch. The switching process of the inverter creates a lot of high-frequency electrical noise. (To produce a sine wave, a special switching scheme is needed to build some kind of approximation to a sinusoid using a series of pulses that then can be filtered to produce a smooth product that looks like normal mains power.)

With a square wave, no such filtering is necessary to produce a power output that the microcomputer will run on. However, the chance that inverter noise will be present in the SPS's output is far greater with a square waveform. Also, a square wave is not a fundamental of 60 Hz as is a sine wave. The "shoulders" of the square waveform contain odd harmonics of the fundamental 60-Hz signal. This means that the manufacturer must take care to eliminate noise from the nonsine wave of an inverter. If this is done, there is no real reason to shy away from nonsine wave products, but true sine wave units typically put out less interference and cost more money.

Synchronizing. When an SPS retransfers to mains power, the waveform output of the inverter needs to be adjusted to match the phase of the incoming utility power. This is called synchronizing or phase matching. When the two waveforms are in phase, no gap will occur when retransfer takes place. A small gap probably will not affect the computer's power supply, but phase matching is a feature that is indicative of good design and possibly makes for more reliable performance.

Low battery warning/shutoff. When batteries power a load during a power outage, their stored energy is slowly depleted. At some point, the depletion is so dramatic that the voltage level of each cell in the battery begins to drop. At a level called end voltage, further discharge will permanently damage the cell. To preserve the life of the battery, most quality SPSs shut off the inverter before this happens. Without this feature, called low battery *shutoff*, the SPS might survive only a few long-term outages.

Ideally, you will shut down the microcomputer and the SPS before this happens. An audible/visible low-battery warning signal is essential to enable the shutdown to be performed in time. More advanced designs can be monitored from a PC, as seen in Figure 5.28 where a piece of Windows software called Check UPS from BEST Power Technology is monitoring one of its UPS units.

Power conditioning. A big sales pitch for UPS makers is that their units also are power conditioners, removing spikes, surges, and noise interference. However, bear in mind that while some systems might do this, there is little inherent protection in most SPS and UPS technology. Indeed, an offline or SPS unit will pass all spikes and surges straight through to your computer unless the design incorporates proper filtering circuits. The switching activity of the inverter in an online system can generate interference over and above that present in the mains supply.

You will recall that the two types of spikes are normal (line/neutral) and common (neutral/ground). The double conversion process of an online UPS will suppress high-energy impulses appearing between line and neutral, but online technology does nothing to prevent common-mode impulses from getting to sensitive equipment. As a matter of fact, all UPS and SPS units generate significant common-mode noise. A transformer can solve this problem because a transformer has its neutral and ground bonded together, thereby shorting out common-mode noise. The transformer itself can be part of an excellent design to thwart common-mode events. The ferroresonant design mentioned earlier is a good power line conditioner for both normal-mode and common-mode events, making the ferroresonant SPS design an increasingly popular choice. However, you can buy separate line conditioners that plug in downstream from an SPS to provide similar benefits.

The network perspective. When most of your personal computers are running in a networked environment, it makes sense to spend most of your power protection budget on the file server. This is where the added cost of a true UPS is justified, because the file server is storing data for more than one user. You still might want to complement this with less expensive standby units on network workstations, printers, external modems, and other peripherals.




Figure 5.28 Check UPS software from BEST Power Technology

Software Assistance

The unexpected loss of power is a hardware problem, but there are some software solutions. The main impact of a power outage is the loss of unsaved work. Many popular personal computer programs do most of their work in RAM, the memory area that receives your input from the keyboard. Your software must arrange for the computer to store your input onto disk; otherwise the input is simply not recorded.

Whatever information is in RAM when the power goes out is wiped out. This includes the operating system and the software application that you were using at the time. You have copies of the application and the operating system on disk. The unsaved input in RAM is what is at risk. You know you are supposed to save your work on a regular basis, but human nature means that this rule will never be adhered to with the diligence required to prevent at least one disaster per user per lifetime.

While you can buy battery-backed RAM, it is more expensive than regular RAM and still quite rare. Most personal computers are simply built with the major weakness of regular RAM as part of the design. To cope with this design weakness, plus human weakness, as well as the fact that UPSs are not universal or cheap, some software provides the user with help in issuing save-my-work commands on a regular basis.

A typical example of a popular application with automatic saving is WordPerfect, which saves the document that you are editing every so many minutes. The time period is up to you and can be adjusted in the software (see Figure 5.29). The actual saving takes only a fraction of a second when you are editing a two- or three-page document. On longer documents or older computers, you might find that you cannot

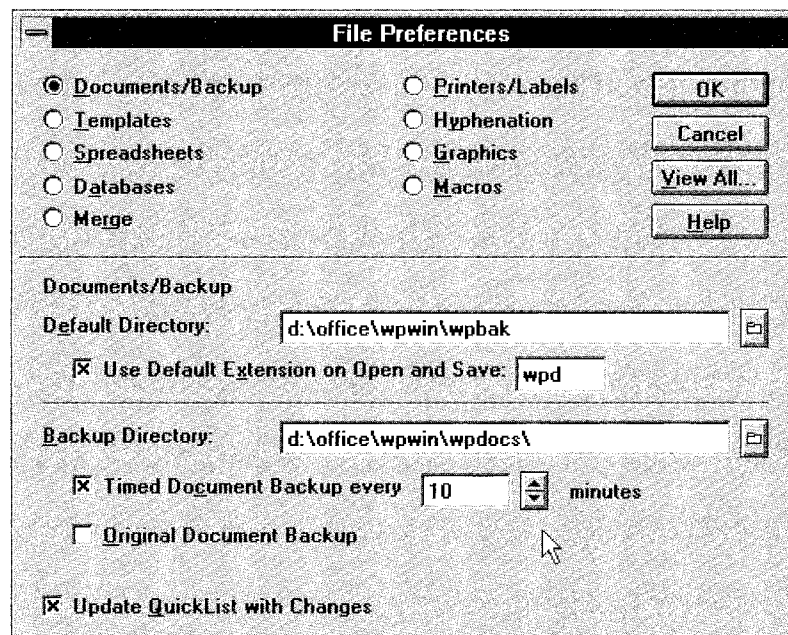


Figure 5.29 Setting the automatic backup time in WordPerfect for Windows

keep typing at high speed throughout the saving process, but it is likely that many automatic saves will occur while you are not actually typing.

The minor inconvenience of a slight delay caused by autosave is vastly outweighed by the added peace of mind that comes from knowing that you will not lose a lot of your editing when the power goes out. Another form of autosave occurs when you use the Stacker disk compression program. This program regularly creates copies of file allocation tables and other information that is vital to recover from a unscheduled system reset or power outage. Yet another type of autosave is complete ongoing system saving, including all applications, so that it is possible to make a full recovery from a crash. Over the years, a number of products have been introduced to address this problem, but none of them has achieved widespread popularity. One reason for this is the widespread use of networks. The file server, which usually is well protected, enables quick recovery on most of the personal computers attached to it.

Electronic Eavesdropping and Worse

Earlier in the chapter, the problems of radio frequency interference (RFI) and electromagnetic interference (EMI) were described from the point of view of interference with the electrical activity within a personal computer. It also was pointed out that personal computers themselves produce RFI and EMI. The high-speed pulsing of electrical current through the circuits of a personal computer gives rise to harmonics—electrical frequencies strong and high enough to be tuned in by radio receivers. This means that activity on a personal computer or terminal can be detected from a distance if you have the right equipment.

The Weak Links

A computer's drives, keyboard, and external connections, like the printer port, emit detectable signals, and it is possible, though challenging, to make sense of these signals. An easier way to learn what a personal computer is doing is to look at the screen. One of the strongest sources of EMI in your personal computer system is the cathode ray tube (CRT) in a typical monitor and the display adapter circuits that drive it. This is because the video signal from the personal computer is boosted to several hundred volts to drive the electron beam of the CRT. The radiated harmonics of the personal computer's video signal are fairly close to those of a typical broadcast TV signal, so all that is required to pick up and read this signal from a distance is a directional antenna and a television. The image will appear scrambled; the signals that provide the horizontal and vertical synchronization do not transmit very well. However, the addition of some fairly cheap circuitry will supply the missing signals, so it is possible to read on a remote screen what is displayed on a personal computer monitor.

In 1985, Wim Van Eck, a Dutch electronics researcher, proved to a number of banks that it was possible, using this technology, to read information from their CRTs at distances of almost a mile away. This gave rise to the term *Van Eck phreaking*. The same principle has been demonstrated by Winn Shwartz, author of *Information Warfare*. While this might sound like an esoteric vulnerability, you must at least

consider defending against it, both to avoid direct losses and also to head off accusations of negligence, should the weakness become more widely exploited. The rapid spread of personal computers has greatly exacerbated the problem because it has significantly increased the number of displays on which sensitive information appears. Furthermore, as in other areas of security, most personal computer designs are woefully lacking in protection features.

A non-CRT display, such as an LCD panel, will reduce, but not eliminate, your exposure to this form of EMI eavesdropping. However, you can take steps to minimize EMI security risks. Obviously, you should minimize the amount of sensitive data that is displayed on the CRT. Such information as account numbers, access codes, and passwords should not be left on the screen and should not be displayed at all if this can be avoided. In chapter 4, you read that, in the interests of simple physical security, computer systems should not be located near windows and doors. This goes double for systems displaying sensitive data. Moving these well within the building will help foil eavesdropping. However, if you have very sensitive data to protect, you might want to contact experts in system shielding and Tempest (see the next section).

Another weak link is the cabling used between personal computers components and between multiple personal computers on a network. A determined eavesdropper could tune in emissions from such cabling and make sense of the data it carried. If the cable is heavy-duty coaxial cable (like the cable used between a TV and a VCR or antenna), then it is probably fairly well insulated. However, the expense of coaxial cable for networks has led many companies to use lightweight cable, called *twisted pair*, that is similar to telephone wiring. This is not as well shielded, so emissions would be easier to detect.

The cables used to connect personal computers, printers, keyboards, and monitors should be the round type, rather than the flat ribbon cables. You should only buy cables that are shielded. Older ribbon cables are not shielded and can give off considerable emissions. A technology that transmits data without EMI is fiberoptic, which uses pulses of light rather than electricity. Fiberoptic cable is currently replacing traditional copper phone wiring in many parts of the world. You can network computers with fiberoptic cabling and thus eliminate the chance that data will be overheard as it travels between computers.

Most personal computer equipment has some form of shielding designed to keep emissions from escaping. The use of a metal casing for the system unit, as in the original IBM PC, is one level of shielding. Plastic-cased components often will have a metal inner shell. Both plastic and metal cases can have a black metallic coating on the inside for further protection. Openings in the casing should be kept to a minimum, which is why blank plates are placed over unoccupied expansion slots (and should be replaced when expansion cards are removed). Openings for ventilation should have metal baffles.

Equipment varies greatly when it comes to the level of shielding because this is an area where manufacturers can cut costs without reducing apparent performance. In the United States, the Federal Communications Commission is charged with regulating emissions from computer equipment and making sure that corners are not cut. The FCC classifies equipment based on how well it performs in this area and has the power to prevent the sale of equipment that gives off too much interference. Many

products are sold with a Class B rating, which means they are known to emit interference but at levels that are "reasonable."

Radiated signals from electronic components can be reflected from flat metal surfaces, actually increasing the power of the signals. When locating your personal computer equipment, avoid placing it on metal desks or work surfaces, and avoid proximity to metal doors, filing cabinets, and dividing screens. On the other hand, a metal box placed around the equipment can prevent signals escaping; the thicker the metal, the greater the protection. Air vents and openings for wires entering the box would need to be baffled. If you are considering going this far, then you might consider the Tempest standard, which is described next.

The Tempest "standard"

During the Cold War, the United States government developed a method of testing computer equipment to measure any emissions that can be used to obtain useful information. Equipment that passes these tests is considered to be secure enough for use in sensitive government applications where the computer data being handled is classified. The tests, or rather the problem they address, are referred to as *Tempest*. The tests themselves are classified. They exist in two forms: one for the United States government (referred to as NACSIM5100A) and another for N.A.T.O. (referred to as AMSG720B).

To get Tempest certification for a product, a manufacturer submits it to the government. If the equipment qualifies, the government then places it on the Tempest Preferred Product List. Contracts to supply the government with computer equipment or to carry out data processing tasks for government agencies can specify Tempest-certified equipment. Nongovernment organizations can purchase equipment that is modified to protect against Tempest-type attacks although they are apparently denied details of government Tempest specifications.

Not surprisingly, modifying a piece of personal computer equipment so that it passes Tempest testing adds considerable cost to the equipment. However, it is possible for anyone to secure their computer activities from electronic prying. Some banks are known to use Tempest-level protection on personal computers involved in sensitive applications. Outside of the United States, it is possible to buy equipment with "Tempest-type" protection. It is worth noting that some countries have laws imposing minimum levels of security on anyone storing data about individuals on computers. For example, the United Kingdom's Data Protection Act extends this protection to basic name and address information. Future legislation along these lines might set minimum levels for security from eavesdropping, placing a burden of protection on those who use computers for database management.

The HERF factor

While some companies have looked into protection against tell-tale waves getting out, very few have taken steps to defend against destructive waves coming in. In *Information Warfare*, Shwartz points out that there is considerable potential for computers and other electronic components to be disrupted by electromagnetic radiation, indeed he opens one chapter with a tantalizing description of how this tech-

nology might be deployed to intercept planes that are smuggling drugs. He uses the term High Energy Radio Frequency gun, or HERF gun, for this technology, which directs high-power radio signals at electronic circuits, causing them to overload. The effect on computers is severe disruption, if not outright hardware damage.

Gentler souls might wonder why on earth an ordinary, law-abiding company should worry about such devices. The answer is a simple, yet chilling phrase: "denial of service." Remember the three pillars of computer security: confidentiality, integrity, and availability. If you have taken steps to protect the confidentiality and integrity of your data with access controls, encryption, backups, and so forth, that still leaves the third pillar: availability. Many organizations now rely on having instant access to their data and constant availability of their processing systems. If this were denied by a burst of high-energy radiation, the costs would be considerable. In other words, if an attacker is not so much interested in grabbing specific data from your systems, but rather is intent on causing general harm, this might be the way to go. If the motive is financial gain, then various extortion scenarios spring to mind like "Give me \$5 million, or your trading system will go down 60 seconds before the New York Stock Exchange opens."

In a military context, the disruptive effect of nuclear explosions, which produce a massive pulse of electromagnetic energy, has long been recognized. The offensive use of this capability, called ElectroMagnetic Pulse Transformer bomb, or EMP/T bomb, was first postulated as a high-altitude nuclear explosion that would destroy all electronic means of communication over a wide area. Indeed, some anti-nuclear groups suggested that this aspect of nuclear bombs alone served to seriously undermine many nuclear warfare scenarios.

After the Cold War, military attention focused on nonnuclear EMP/T devices for use in limited theaters of operation. However, the offensive potential of HERF guns and EMP/T bombs in domestic scenarios, such as extortion or terrorism, has largely been ignored, and the subject remains on the fringes of mainstream computer security. Given that no attacks of this type have yet been publicly reported, this is where we must leave the subject, but my advice is to "stay tuned" for any new developments. If attacks of this type do materialize, then a whole new range of defensive measures will need to be considered.

The Network Connection

Clearly, the network file server should be equipped with surge protection and an intelligent backup power supply that can execute, through software, an orderly shutdown of the network if there is a prolonged power outage. Typically, this will involve broadcasting an e-mail warning to all network stations and logging them off the server. Fortunately, such technology is now readily available and competitively priced.

However, there is not much point having a file server that can notify network workstations of an imminent shutdown due to power outage, if they have already lost power. The answer is UPS protection for all computers on the network, which now is available for less than \$100 per machine. Models such as the ViewSonic Opti-UPS 280E have set a new price point, offering 280VA backup, which is more than enough


for most desktops machines, using a "line-interactive" design. While not fully online, this new type of UPS keeps the inverter running to top up the battery and simply reverses it when there is a mains outage, resulting in a very rapid response time. This design also provides "buck-boost" protection to smooth out both under- and over-voltage conditions, without tapping battery power.

Summary

One of the main reasons that the use of personal computers has spread so far so fast is the low cost of hardware and software relative to mainframe and minicomputers. In part, this low cost has been achieved by simply ignoring some aspects of design that are considered essential in larger systems. Power conditioning and backup are either built into mainframe computers or installed as a matter of course when the system is first set up. Many mainframes have disaster recovery routines built into their operating system software or provided for in their application programs.

As personal computers take on more and more of the tasks previously performed by larger systems, the true cost of personal computer systems becomes increasingly apparent. Serious applications require disaster recovery features. Most installations need some form of power conditioning. All vital systems should have power backup so that they can keep computing during brief power failures and accomplish an orderly shutdown during prolonged outages. While provision of these features increases the price tag for personal computer technology, most users will consider the money well spent, providing a valuable security against data loss, hardware damage, and costly interruptions in normal operations.

"Used by kind permission of Paul Davies Publications"



Expect
the
unexpected
—
Back up
now!