Chapter 18

Maintaining Your System: Preventive Maintenance, Backups, and Warranties

Preventive maintenance is the key to obtaining years of trouble-free service from your computer system. A properly administered preventive maintenance program pays for itself by reducing problem behavior, data loss, component failure, and by ensuring a long life for your system. In several cases, I have "repaired" an ailing system with nothing more than a preventive maintenance session. Preventive maintenance also increases your system's resale value because it will look and run better. This chapter describes preventive maintenance procedures and how often you should perform them.

You will also learn the importance of creating backup files of data and the various backup procedures available. A sad reality in the computer repair and servicing world is that hardware can always be repaired or replaced, but data cannot. Most hard disk troubleshooting and service procedures, for example, require that a low-level format be done. This low-level format overwrites any data on the disk.

Because data recovery depends a great deal on the type and severity of damage and the expertise of the recovery specialist, data-recovery services are very expensive. Most recovery services charge a premium and offer no guarantees that the data will be completely recovered. Backing up your system as discussed in this chapter is the only guarantee you have of seeing your data again.

Most of the discussion of backing up systems in this chapter is limited to professional solutions that require special hardware and software. Backup solutions that employ floppy disk drives, such as the DOS backup software, are insufficient and too costly in most cases for hard disk backups. It would take 1,456 1.44M floppy disks, for example, to backup the 2G hard disk in my portable system! That would cost more than $1,000 in disks, not to mention the time involved. A tape system can put 4G to 8G on a single $15 tape.

Finally, the last section in this chapter discusses the standard warranties and optional service contracts available for many systems. Although most of this book is written for people who want to perform their own maintenance and repair service, taking advantage of a good factory warranty that provides service for free definitely is prudent. Some larger computer companies, such as IBM, offer attractive service contracts that, in some cases, are cost-justified over self service. These types of options are examined in the final section.

Developing a Preventive Maintenance Program

Developing a preventive maintenance program is important to everyone who uses or manages personal computer systems. Two types of preventive maintenance procedures exist: active and passive.

Active preventive maintenance includes steps you apply to a system that promote a longer, trouble-free life. This type of preventive maintenance primarily involves periodic cleaning of the system and its components. This section describes several active preventive maintenance procedures, including cleaning and lubricating all major components, reseating chips and connectors, and reformatting hard disks.

Passive preventive maintenance includes steps you can take to protect a system from the environment, such as using power-protection devices; ensuring a clean, temperature-controlled environment; and preventing excessive vibration. In other words, passive preventive maintenance means treating your system well. This section also describes passive preventive maintenance procedures.

Active Preventive Maintenance Procedures

How often you should implement active preventive maintenance procedures depends on the system's environment and the quality of the system's components. If your system is in a dirty environment, such as a machine shop floor or a gas station service area, you might need to clean your system every three months or less. For normal office environments, cleaning a system every one to two years is usually fine. However, if you open your system after one year and find dust bunnies inside, you should shorten the cleaning interval.

Another active preventive maintenance technique discussed in this section is reformatting hard disks. Low-level reformatting restores the track and sector marks to their proper locations and forces you to back up and restore all data on the drive. Not all drives require this procedure, but if you are using drives with a stepper motor head actuator, periodic reformatting is highly recommended. Most drives with voice-coil head actuators run indefinitely without reformatting due to their track following servo mechanisms which prevent temperature induced mistracking.

Other hard disk preventive maintenance procedures include making periodic backups of critical areas such as Boot Sectors, File Allocation Tables, and Directory structures on the disk.

Cleaning a System. One of the most important operations in a good preventive maintenance program is regular and thorough cleaning of the system. Dust buildup on the internal components can lead to several problems. One is that dust acts as a thermal insulator which prevents proper system cooling. Excessive heat shortens the life of system components and adds to the thermal stress problem caused by wider than normal temperature changes between power-on and power-off states. Additionally, dust may contain conductive elements that can cause partial short circuits in a system. Other elements in dust and dirt can accelerate corrosion of electrical contacts and cause improper connections. In all, the removal of any layer of dust and debris from within a computer system benefits that system in the long run.

All IBM and IBM-compatible systems use a forced-air cooling system that allows for even cooling inside the system. A fan is mounted in, on, or near the power supply and pushes air outside. This setup depressurizes the interior of the system relative to the outside air. The lower pressure inside the system causes outside air to be drawn into openings in the system chassis and cover. This draw-through, or depressurization, is the most efficient cooling system that can be designed without an air filter. Air filters typically are not used with depressurization systems because there is no easy way to limit air intake to a single port that can be covered by a filter.

Some industrial computers from IBM and other companies use a forced-air system that uses the fan to pressurize, rather than to depressurize, the case. This system forces air to exhaust from any holes in the chassis and case or cover. The key to the pressurization system is that all air intake for the system is at a single location—the fan. The air flowing into the system therefore can be filtered by simply integrating a filter assembly into the fan housing. The filter must be cleaned or changed periodically. Because the interior of the case is pressurized relative to the outside air, airborne contaminants are not drawn into the system even though it may not be sealed. Any air entering the system must pass through the fan and filter housing, which removes the contaminants. Pressurization cooling systems are used primarily in industrial computer models designed for extremely harsh environments.

Most systems you have contact with are depressurization systems. Mounting any sort of air filter on these types of systems is impossible because air enters the system from too many sources. With any cooling system in which incoming air is not filtered, dust and other chemical matter in the environment is drawn in and builds up inside the computer. This buildup can cause severe problems if left unchecked.

One problem that can develop is overheating. The buildup of dust acts as a heat insulator, which prevents the system from cooling properly. Some of the components in a modern PC can generate an enormous amount of heat which must be dissipated for the component to function. The dust also might contain chemicals that conduct electricity. These chemicals can cause minor current shorts and create electrical signal paths where none should exist. The chemicals also cause rapid corrosion of cable connectors, socket-installed components, and areas where boards plug into slots. All can cause intermittent system problems and erratic operation.

Tip

Cigarette smoke contains chemicals that can conduct electricity and cause corrosion of computer parts. The smoke residue can infiltrate the entire system, causing corrosion and contamination of electrical contacts and sensitive components such as floppy drive read/write heads and optical drive lens assemblies. You should avoid smoking near computer equipment and encourage your company to develop and enforce a similar policy.

Floppy disk drives are particularly vulnerable to the effects of dirt and dust. Floppy drives are a large "hole" within the system through which air continuously flows. Therefore, they accumulate a large amount of dust and chemical buildup within a short time. Hard disk drives do not present quite the same problem. Because the head disk assembly (HDA) in a hard disk is a sealed unit with a single barometric vent, no dust or dirt can enter without passing through the barometric vent filter. This filter ensures that contaminating dust or particles cannot enter the interior of the HDA. Thus, cleaning a hard disk requires simply blowing the dust and dirt off from outside the drive. No internal cleaning is required.

Disassembly and Cleaning Tools. To properly clean the system and all the boards inside requires certain supplies and tools. In addition to the tools required to disassemble the unit (see Chapter 3, "System Teardown and Inspection"), you should have these items:

- Contact cleaning solution
- Canned air
- A small brush
- Lint-free foam cleaning swabs
- Antistatic wrist-grounding strap

You might also want to acquire these optional items:

- Foam tape
- Low-volatile room-temperature vulcanizing (RTV) sealer
- Silicone type lubricant
- Computer vacuum cleaner

These simple cleaning tools and chemical solutions will allow you to perform most common preventive maintenance tasks.

Chemicals. You can use several different types of cleaning solutions with computers and electronic assemblies. Most fall into the following categories:

- Standard Cleaners
- Contact Cleaner/Lubricants
- Dusters

Tip

The makeup of many of the chemicals used for cleaning electronic components has been changing because many of the chemicals originally used are now considered environmentally unsafe. They have been attributed to damaging the Earth's ozone layer, a natural protective barrier in the stratosphere which prevents harmful ultraviolet (UV-B) radiation from reaching Earth. Chlorine atoms from chlorofluorocarbons (CFCs) and chlorinated solvents attach themselves to ozone molecules and destroy them. Many of these chemicals are now strictly regulated by federal and international agencies in an effort to preserve the ozone layer. Most of the companies that produce chemicals used for system cleaning and maintenance have had to introduce environmentally safe replacements. The only drawback is that many of these safer chemicals cost more and usually do not work as well as those they replace.

Many specific chemicals are used in cleaning and dusting solutions, but five types are of particular interest. The EPA has classified ozone-damaging chemicals into two classes: Class I and Class II. Chemicals that fall into these two classes have their usage regulated. Other chemicals are nonregulated. Class I chemicals include:

- Chlorofluorocarbons (CFCs)
- Chlorinated solvents

Class I chemicals can only be sold for use in professional service and not to consumers. A new law that went into effect on May 15, 1993, requires that the containers for Class I chemicals be labeled with a warning that the product "Contains substances which harm public health and the environment by destroying ozone in the atmosphere." Additionally, electronics manufacturers and other industries must also apply a similar warning label to any products that use Class I chemicals in the production process. This means that any circuit board or computer that is manufactured with CFCs will have this label!

The most popular Class I chemicals are the various forms of Freon, which are CFCs. A very popular cleaning solution called 1,1,1 Trichloroethane is a chlorinated solvent and is also strictly regulated. Up until the last year or so, virtually all computer or electronic cleaning solutions contained one or both of these chemicals. While you can still purchase them, regulations and limited production have made them more expensive and more difficult to find.

Class II chemicals include hydrochlorofluorocarbons (HCFCs). These are not as strictly regulated as Class I chemicals because they have a lower ozone depletion potential. Many cleaning solutions have switched to HCFCs, because they do not require the restrictive labeling required by Class I chemicals and are not as harmful. Most HCFCs have only one tenth the ozone damaging potential of CFCs.

Other nonregulated chemicals include Volatile Organic Compounds (VOCs) and Hydrofluorocarbons (HFCs). These chemicals do not damage the ozone layer but actually contribute to ozone production, which, unfortunately, appears in the form of smog or ground level pollution. Pure isopropyl alcohol is an example of a VOC that is commonly used in electronic part and contact cleaning. HFCs are used as a replacement for CFCs because the HFCs do not damage the ozone layer.

The EPA has developed a method to measure the ozone damaging capability of a chemical. The Ozone Depletion Potential (ODP) of a chemical solution is the sum of the depletion potentials of each of the chemicals used in the solution by weight. The ODP of Freon R12 (Automotive Air Conditioning Freon), is 1.0 on this scale. Most modern CFC replacement chemicals have an ODP rating of 0.0 to 0.1, as opposed to those using CFCs and chlorinated solvents that usually have ODP ratings of 0.75 or higher.

Standard Cleaners. Standard cleaning solutions are available in a variety of types and configurations. You can use pure isopropyl alcohol, acetone, Freon, trichloroethane, or a variety of other chemicals. Most are now leaning to the alcohol, acetone, or others that do not cause ozone depletion and that comply with government regulations and environmental safety. You should be sure that your cleaning solution is designed to clean computers or electronic assemblies. In most cases this means that the solution should be chemically pure and free from contaminants or other unwanted substances. You should not, for example, use drugstore rubbing alcohol for cleaning electronic parts or contacts because it is not pure and could contain water or perfumes. The material must be moisture-free and residue-free. The solutions should be in liquid form, not a spray. Sprays can be wasteful and you almost never spray the solution directly on components. Instead, wet a foam or chamois swab used for wiping the component. These electronic-component cleaning solutions are available at any good electronics parts stores.

Contact Cleaner/Lubricants. These are very similar to the standard cleaners but include a lubricating component. The lubricant eases the force required when plugging and unplugging cables and connectors, which reduces strain on the devices. The lubricant coating also acts as a conductive protectant that insulates the contacts from corrosion. These chemicals can greatly prolong the life of a system by preventing intermittent contacts in the future.

Contact cleaner/lubricants are especially effective on I/O slot connectors, adapter card edge and pin connectors, disk drive connectors, power supply connectors, and virtually all connectors in the PC.

An excellent contact enhancer and lubricant is Stabilant 22. It is more effective than conventional contact cleaners or lubricants. This chemical is available in several forms. Stabilant 22 is the full strength concentrated version, while Stabilant 22a is a version diluted with isopropyl alcohol in a 4 to 1 ratio. An even more diluted 8 to 1 ratio version is sold in many high end stereo and audio shops under the name "Tweek." Just 15ml of Stabilant 22a sells for about $40, while a liter of the concentrate costs about $4,000. While Stabilant 22 is expensive, very little is required and an application can provide protection for a long time. Stabilant is manufactured by D. W. Electrochemicals, which is listed in Appendix B, the "Vendor List."

Dusters. Compressed gas often is used as an aid in system cleaning. The compressed gas is used as a blower to remove dust and debris from a system or component. Originally, these dusters used CFCs such as Freon, while modern dusters now use HFCs or carbon dioxide, neither of which is damaging to the ozone layer. Be careful when you use these devices because some of them can generate a static charge when the compressed gas leaves the nozzle of the can. Be sure that you are using the kind approved for cleaning or dusting off computer equipment, and consider wearing a static grounding strap as a precaution. The type of compressed-air cans used for cleaning camera equipment can sometimes differ from the type used for cleaning static-sensitive computer components.

Most older computer-grade canned gas dusters consisted of dichlorodifluoromethane (Freon R12), the same chemical used in many automotive air-conditioning systems built until 1995, when a ban of the manufacture and use of R12 takes place. In 1992 many automobile manufacturers began switching to an ozone safe chemical called R134a. Manufacturing of R12 will cease by 1995, and the regulations placed on its use have forced companies to use other products such as carbon dioxide for compressed gas dusters. In addition to the environmental concerns about depleting the ozone layer, Freon can be dangerous if exposed to an open flame.

Caution

If you use any chemical with the propellant Freon R12 (dichlorodifluoromethane), *do not expose the gas to an open flame or other heat source.* If you burn this substance, a highly toxic gas called *phosgene* is generated. Phosgene, used as a nerve gas in World War II, can be deadly.

Related to compressed-air products are chemical-freeze sprays. These sprays are used to quickly cool down a suspected failing component, which often temporarily restores it to operation. These substances are not used to repair a device, but to confirm that you have found a failed device. Often, a component's failure is heat-related and cooling it temporarily restores it to function. If the circuit begins operating normally, the device you are cooling is the suspect device.

Vacuum Cleaners. Some people prefer to use a vacuum cleaner instead of canned gas dusters for cleaning a system. Canned gas is usually better for cleaning in small areas. A vacuum cleaner is more useful when you are cleaning a system loaded with dust and dirt. You can use the vacuum cleaner to suck out dust and debris instead of blowing them on other components, which sometimes happens with canned air. For outbound servicing (when you are going to the location of the equipment instead of the equipment coming to you), canned air is easier to carry in a toolkit than a small vacuum cleaner.

Brushes and Swabs. A small brush (makeup, photographic, or paint) can be used to carefully loosen accumulated dirt and dust before spraying with canned air or using the vacuum cleaner. Be careful about generating static electricity. In most cases the brushes should not be used directly on circuit boards but should be used instead on the case interior and other parts such as fan blades, air vents, and keyboards. Wear a grounded wrist strap if you are brushing on or near any circuit boards, and brush slowly and lightly to prevent static discharges from occurring.

Use cleaning swabs to wipe off electrical contacts and connectors, disk drive heads, and other sensitive areas. The swabs should be made of foam or synthetic chamois material that does not leave lint or dust residue. Unfortunately, proper foam or chamois cleaning swabs are more expensive than the typical cotton swabs. Do not use cotton swabs because they leave cotton fibers on everything they touch. Cotton fibers are conductive in some situations and can remain on drive heads, which can scratch disks. Foam or chamois swabs can be purchased at most electronics-supply stores.

One item to avoid is an eraser for cleaning contacts. Many people (including myself) have recommended using a soft pencil type eraser for cleaning circuit board contacts. Testing has proven this to be bad advice for several reasons. One is that any such abrasive wiping on electrical contacts generates friction and an electrostatic discharge (ESD). This ESD can be damaging to boards and components, especially with newer low voltage devices made using CMOS (Complimentary Metal Oxide Semiconductor) technology. These devices are especially static sensitive, and cleaning the contacts without a proper liquid solution is not recommended. Also the eraser will wear off the gold coating on many contacts, exposing the tin contact underneath, which will rapidly corrode when exposed to air. Some companies sell premoistened contact cleaning pads that are soaked in a proper contact cleaner and lubricant. These pads are safe to wipe on conductor and contacts with no likelihood of ESD damage or abrasion of the gold plating.

Foam Tape or RTV Sealer. Hard disks usually use a small copper strap to ground the spindle of the disk assembly to the logic board, thus bleeding off any static charge carried by the spinning disk platters. Unfortunately, this strap often can begin to harmonize, or vibrate, and result in an annoying squealing or whining noise. (Sometimes the noise is similar to fingernails dragged across a chalkboard.)

To eliminate the source of irritation, you can stop the strap from vibrating by weighting it. One method is to use a piece of foam tape cut to match the size of the strap and stuck to the strap's back side. Another way to dampen the vibration is to apply a low-volatile RTV sealer. You apply this silicone-type rubber to the back of the grounding strap. After it hardens to a rubber-like material, the sealer stops the vibrations that produce the annoying squeal. You can buy the RTV sealer from an automotive-supply house.

I prefer using the foam tape rather than the RTV because it is easier and neater to apply. If you use the RTV, be sure that it is the low-volatile type, which does not generate acid when it cures. This acid produces the vinegar smell common to the standard RTV sealer, and can be highly corrosive to the strap and anything else it contacts. The low-volatile RTV also eliminates the bad vinegar smell. You can purchase the foam tape at most electronics-supply houses, where it often is sold for attaching alarm switches to doors or windows. The low-volatile RTV is available from most auto-supply stores. To be sure that you buy low-volatile RTV, look for the packaging to state specifically that the product is either a low-volatile type or is compatible with automobile oxygen sensors.

Silicone Lubricants. Silicone lubricants are used to lubricate the door mechanisms on floppy disk drives and any other part of the system that may require clean, non-oily lubrication. Other items you can lubricate are the disk drive head slider rails or even printer-head slider rails, which allow for smooth operation.

Using silicone instead of conventional oils is important because silicone does not gum up and collect dust and other debris. Always use the silicone sparingly. Do not spray it anywhere near the equipment as it tends to migrate and will end up where it doesn't belong (such as on drive heads). Instead, apply a small amount to a toothpick or foam swab and dab the silicone on the components where needed. You can use a lint-free cleaning stick soaked in silicone lubricant to lubricate the metal print-head rails in a printer.

Remember that some of the cleaning operations described in this section might generate a static charge. You may want to use a static grounding strap in cases in which static levels are high to ensure that you do not damage any boards as you work with them.

Obtaining Required Tools and Accessories. Most cleaning chemicals and tools can be obtained from a number of electronics-supply houses, or even the local Radio Shack. A company called Chemtronics specializes in chemicals for the computer and electronics industry. These and other companies which supply tools, chemicals, and other computer and electronic cleaning supplies are listed in Appendix B, the "Vendor List." With all these items on hand, you should be equipped for most preventive maintenance operations.

Disassembling and Cleaning Procedures. To properly clean your system, it must at least be partially disassembled. Some people go as far as to remove the motherboard. Removing the motherboard results in the best possible access to other areas of the system but in the interest of saving time, you probably need to disassemble the system only to where the motherboard is completely visible.

All plug-in adapter cards must be removed, along with the disk drives. Although you can clean the heads of a floppy drive with a cleaning disk without opening the system unit's cover, you probably will want to do more thorough cleaning. In addition to the heads, you also should clean and lubricate the door mechanism and clean any logic boards and connectors on the drive. This procedure usually requires removing the drive.

Next, do the same procedure with a hard disk: clean the logic boards and connectors, as well as lubricate the grounding strap. To do so, you must remove the hard disk assembly. As a precaution, be sure it is backed up before removal.

Reseating Socketed Chips. A primary preventive maintenance function is to undo the effects of chip creep. As your system heats and cools, it expands and contracts, and the physical expansion and contraction causes components plugged into sockets to gradually work their way out of those sockets. This process is called *chip creep*. To correct its effects, you must find all socketed components in the system and make sure that they are properly reseated.

In most systems, all the memory chips are socketed or are installed in socketed SIMMs (Single Inline Memory Modules). SIMM devices are retained securely in their sockets by a positive latching mechanism and cannot creep out. Memory SIPP (Single Inline Pin Package) devices (SIMMs with pins rather than contacts) are not retained by a latching mechanism and therefore can creep out of their sockets. Standard socketed memory chips are prime candidates for chip creep. Most other logic components are soldered in. You can also expect to find the ROM chips, the main processor or CPU, and the math coprocessor in sockets. In most systems, these items are the only components that are socketed; all others are soldered in.

Exceptions, however, might exist. A socketed component in one system might not be socketed in another—even if both are from the same manufacturer. Sometimes this difference results from a parts-availability problem when the boards are manufactured. Rather than halt the assembly line when a part is not available, the manufacturer adds a socket instead of the component. When the component becomes available, it is plugged in and the board is finished. Many newer systems place the CPU in a Zero Insertion Force (ZIF) socket, which has a lever that can release the grip of the socket on the chip. In most cases there is very little creep with a ZIF socket.

To make sure that all components are fully seated in their sockets, place your hand on the underside of the board and then apply downward pressure with your thumb (from the top) on the chip to be seated. For larger chips, seat the chip carefully in two movements, and press separately on each end of the chip with your thumb to be sure that the chip is fully seated. (The processor and math coprocessor chips can usually be seated in this manner.) In most cases, you hear a crunching sound as the chip makes its way back into the socket. Because of the great force sometimes required to reseat the chips, this operation is difficult if you do not remove the board.

For motherboards, forcibly seating chips can be dangerous if you do not directly support the board from the underside with your hand. Too much pressure on the board can cause it to bow or bend in the chassis, and the pressure can crack it before seating takes place. The plastic standoffs that separate and hold the board up from the metal chassis are spaced too far apart to properly support the board under this kind of stress. Try this operation only if you can remove and support the board adequately from underneath.

You may be surprised to know that, even if you fully seat each chip, they might need reseating again within a year. The creep usually is noticeable within a year or less.

Cleaning Boards. After reseating any socketed devices that may have crept out of their sockets, the next step is to clean the boards and all connectors in the system. For this step, the cleaning solutions and the lint-free swabs described earlier are needed.

First, clean the dust and debris off the board and then clean any connectors on the board. To clean the boards, it is usually best to use a vacuum cleaner designed for electronic assemblies and circuit boards or a duster can of compressed gas. The dusters are especially effective at blasting any dust and dirt off the boards.

Also blow any dust out of the power supply, especially around the fan intake and exhaust areas. You do not need to disassemble the power supply to do this, just use a duster can and blast the compressed air into the supply through the fan exhaust port. This will blow the dust out of the supply and clean off the fan blades and grille which will help with system airflow.

Caution

Be careful with electrostatic discharge (ESD), which can damage components, when cleaning electronic components. Take extra precautions in the dead of winter in an extremely dry, high-static environment. You can apply antistatic sprays and treatments to the work area to reduce the likelihood of ESD damage.

An antistatic wrist-grounding strap is recommended. This should be connected to a ground on the card or board you are wiping. This strap ensures that no electrical discharge occurs between you and the board. An alternative method is to keep a finger or thumb on the ground of the motherboard or card as you wipe it off. It is easier to ensure proper grounding while the motherboard is still installed in the chassis, so it is a good idea not to remove it.

Cleaning Connectors and Contacts. Cleaning the connectors and contacts in a system promotes reliable connections between devices. On a motherboard, you will want to clean the slot connectors, power-supply connectors, keyboard connector, and speaker connector. For most plug-in cards, you will want to clean the edge connectors that plug into slots on the motherboard as well as any other connectors, such as external ones mounted on the card bracket.

Submerge the lint-free swabs in the liquid cleaning solution. If you are using the spray, hold the swab away from the system and spray a small amount on the foam end until the solution starts to drip. Then, use the soaked foam swab to wipe the connectors on the boards. Pre-soaked wipes are the easiest to use. Simply wipe them along the contacts to remove any accumulated dirt and leave a protective coating behind.

On the motherboard, pay special attention to the slot connectors. Be liberal with the liquid; resoak the foam swab repeatedly, and vigorously clean the connectors. Don’t worry if some of the liquid drips on the surface of the motherboard. These solutions are entirely safe for the whole board and will not damage the components.

Use the solution to wash the dirt off the gold contacts in the slot connectors, and then douse any other connectors on the board. Clean the keyboard connector, the grounding positions where screws ground the board to the system chassis, power-supply connectors, speaker connectors, battery connectors, and so on.

If you are cleaning a plug-in board, pay special attention to the edge connector that mates with the slot connector on the motherboard. When people handle plug-in cards, they often touch the gold contacts on these connectors. Touching the gold contacts coats them with oils and debris, which prevents proper contact with the slot connector when the board is installed. Make sure that these gold contacts are free of all finger oils and residue. It is a good idea to use one of the contact cleaners that has a conductive lubricant, which both allows connections to be made with less force, and also protects the contacts from corrosion.

Caution

Many people use a common pink eraser to rub the edge connectors clean. I do not recommend this procedure for two reasons. One, the eraser eventually removes some of the gold and leaves the tin solder or copper underneath exposed. Without the gold, the contact corrodes rapidly and requires frequent cleaning. The second reason to avoid cleaning with the eraser is that the rubbing action can generate a static charge. This charge can harm any component on the board. Rather than use an eraser, use the liquid solution and swab method described earlier.

You also will want to use the swab and solution to clean the ends of ribbon cables or other types of cables or connectors in a system. Clean the floppy drive cables and connectors, the hard disk cables and connectors, and any others you find. Don't forget to clean off the edge connectors that are on the disk drive logic boards as well as the power connectors to the drives.

Cleaning Floppy Drives. Because Chapter 13, "Floppy Disk Drives and Controllers," explains the procedure for cleaning floppy drives, the information is not repeated here. The basic idea is to use a canned gas duster to dust off the interior of the drive, use the silicone lubricant on whatever items need lubrication, and follow with a head cleaning, either manually with a foam swab or most likely with a chemical-soaked cleaning disk.

For hard disks, take this opportunity to dampen or lubricate the grounding strap if you have a noise problem as described earlier. Dampening is the recommended solution because if you lubricate this point, the lubricant eventually dries up and the squeal can come back. Because the dampening is usually a more permanent fix for this sort of problem, I recommend it whenever possible. Most newer hard disks have this dampening material applied at the factory and are not likely to generate noise like older drives.

Cleaning the Keyboard and Mouse. Keyboards and mice are notorious for picking up dirt and garbage. If you ever open an older keyboard, you often will be amazed at the junk you will find.

To prevent problems, it is a good idea to periodically clean out the keyboard with a vacuum cleaner. An alternative method is to turn the keyboard upside down and shoot it with a can of compressed gas. This will blow out the dirt and debris that has accumulated inside the keyboard and possibly prevent future problems with sticking keys or dirty keyswitches.

If a particular key is stuck or making intermittent contact, you can soak or spray that switch with contact cleaner. The best way to do this is to first remove the keycap and then spray the cleaner into the switch. This usually does not require complete disassembly of the keyboard. Periodic vacuuming or compressed gas cleaning will prevent more serious problems with sticking keys and keyswitches.

Most mice are easily cleaned. In most cases there is a twist off locking retainer that keeps the mouse ball retained in the body of the mouse. By removing the retainer, the ball will drop out. After removing the ball, you should clean it with one of the electronic cleaners. I would recommend a pure cleaner instead of a contact cleaner with lubricant because you do not want any lubricant on the mouse ball. Then you should clean off the rollers in the body of the mouse with the cleaner and some swabs.

Periodic cleaning of a mouse in this manner will eliminate or prevent skipping or erratic movement that can be frustrating. I also recommend a mouse pad for most ball type mice because the pad will prevent the mouse ball from picking up debris from your desk.

Mice often need frequent cleaning before they start sticking and jumping, which can be frustrating. If you never want to clean a mouse again, I suggest you look into the Honeywell mouse. These mice have a revolutionary new design that uses two external wheels rather than the conventional ball and roller system. The wheels work directly on the desk surface and are unaffected by dirt and dust. Because the body of the mouse is sealed, dirt and dust cannot enter it and gum up the positional sensors. I find this mouse excellent to use with my portable system because it works well on any surface. This mouse is virtually immune to the sticking and jumping that plague ball and roller designs and never needs to be cleaned, so it is less frustrating than conventional mice.

Hard Disk Maintenance

Certain preventive maintenance procedures protect your data and ensure that your hard disk works efficiently. Some of these procedures actually minimize wear and tear on your drive, which will prolong its life. Additionally, a high level of data protection can be implemented by performing some simple commands periodically. These commands provide methods for backing up (and possibly later restoring) critical areas of the hard disk that, if damaged, would disable access to all your files.

Defragmenting Files. Over time, as you delete and save files to a hard disk, the files become fragmented. This means that they are split into many noncontiguous areas on the disk. One of the best ways to protect both your hard disk and the data on it is to periodically defragment the files on the disk. This serves two purposes. One is that by ensuring that all of the files are stored in contiguous sectors on the disk, head movement and drive wear and tear will be minimized. This has the added benefit of improving the speed at which files will be retrieved from the drive by reducing the head thrashing that occurs every time a fragmented file is accessed.

The second major benefit, and in my estimation the more important of the two, is that in the case of a disaster where the File Allocation Tables (FATs) and Root Directory are severely damaged, the data on the drive can usually be recovered very easily if the files are contiguous. On the other hand, if the files are split up in many pieces across the drive, it is virtually impossible to figure out which pieces belong to which files without an intact File Allocation Table (FAT) and directory system. For the purposes of data integrity and protection, I recommend defragmenting your hard disk drives on a weekly basis, or immediately after you perform any major backup.

Three main functions are found in most defragmenting programs:

- File Defragmentation
- File Packing (Free Space Consolidation)
- File Sorting

Defragmentation is the basic function but most other programs also add file packing. Packing the files is optional on some programs because it usually takes additional time to perform. This function packs the files at the beginning of the disk so that all free space is consolidated at the end of the disk. This feature minimizes future file fragmentation by eliminating any empty holes on the disk. Because all free space is consolidated into one large area, any new files written to the disk will be able to be written in a contiguous manner with no fragmentation necessary.

The last function, file sorting, is not usually necessary and is performed as an option by many defragmenting programs. This function adds a tremendous amount of time to the operation, and has little or no effect on the speed at which information is accessed. It can be somewhat beneficial for disaster recovery purposes because you will have an idea of which files came before or after other files if a disaster occurs. These benefits are minimal compared to having the files be contiguous no matter what their order. Not all defragmenting programs offer file sorting, and the extra time it takes is probably not worth any benefits you will receive. Other programs can sort the order that files are listed in directories, which is a quick and easy operation compared to sorting the file ordering the disk.

Several programs are available that can defragment the files on a hard disk but the one available to most people is the DEFRAG program built into DOS 6.x. This reduced function version of the SPEEDISK program is a part of the Symantec Norton Utilities. If you have the Norton Utilities, by all means use SPEEDISK instead of DEFRAG. The DOS DEFRAG program offers all three functions including defragmenting, packing, and sorting. The SPEEDISK program performs these operations faster and with more efficient use of memory. Many of these programs will have problems on very large disks, and SPEEDISK will work on drives as large as 2G (the maximum DOS volume size), while DEFRAG is limited to disks of about 512M or less due to memory constraints.

Several aftermarket defragmenting programs are more powerful or faster than the DOS DEFRAG program. They include:

- SPEEDISK by Symantec (Norton Utilities)
- Power Disk by PC-KWIK
- Optune by Gazelle
- VOPT by Golden Bow

All of these are highly recommended and usually perform much better than the DEFRAG program in DOS. The Power Disk, Optune, and VOPT programs are also much faster than

DEFRAG. The VOPT program is the simplest and quickest of these. It is a command line program not a menu driven program like the others and does not offer a sort capability while all the others do. What it does offer is the ultimate in speed and efficiency in defragmenting and packing files. Currently, most of these programs will work on drives up to 1G while some will handle drives of up to 2G. These programs are updated constantly; contact the manufacturer for more detailed specifications if you are interested. Each of these manufacturers is listed in Appendix B, the "Vendor List."

No matter which program you end up using, defragmenting and packing your disk helps to reduce drive wear and tear by minimizing the amount of work required to load files. It also greatly increases the chances for data recovery in the case of serious corruption in the File Allocation Tables (FATs) and directories on the disk.

Backing Up the FAT and Directory System. The operating system uses several areas on a formatted disk to manage the files stored on the disk. These areas are extremely critical. If they are damaged, all access to the drive volume may be compromised or completely disabled. In some cases, these critical areas can be rebuilt using known data recovery procedures and tools, but the easiest and best way to recover from damage to these areas is to simply have a backup of them to restore.

The critical areas of a hard disk file system are the following:

- Master (Partition) Boot Record (MBR)
- Extended (Partition) Boot Record (EBR)
- DOS Boot Record (DBR)
- File Allocation Tables (FATs)
- Root Directory

These areas are stored on the disk in the order listed, except that each volume starts with an MBR or EBR but not both. Unlike complete file backups of the entire hard disk, a backup of these system areas is relatively quick and easy. This is because these areas are very small in comparison to the remainder of the drive and normally occupy a fixed amount of space on the disk. For example, the MBR, EBRs and DBRs are only one sector long each, while the two FATs used on each volume cannot exceed 512 sectors (256K), and the Root Directory is limited to 32 sectors (15K). This means that even for the largest hard disk drive, these areas do not consume more than about 300K of space.

The boot record (sector) areas do not change during day to day usage of the disk. These areas will change only if you reformat, repartition, or change operating system versions. Because they are relatively constant, it makes good sense to back up these areas to a file on a floppy disk for later, to restore if necessary.

Each normal disk volume will have two File Allocation Tables and a single Root Directory. These areas change constantly as files are written and deleted from the disk. A backup of these areas is only good for temporary purposes and can be useful when trying to undelete files. Because of this, these backups are often written on the hard disk in a special hidden file near the end of the disk. Other recovery or restoration programs can then look for this special file and use the information within it to rebuild the FATs or directory system. Often this type of backup is very useful in a situation where files need to be undeleted.

DOS 5.0 provided the MIRROR command for backups of all of these areas. MIRROR has two functions: it can back up the boot sector areas on the hard disk to a file on a floppy disk as well as back up the FAT and directory areas to the end of the hard disk as a special hidden file that recovery programs can find. To use MIRROR to back up the boot sectors to a floppy disk, execute the command as follows:

```
MIRROR /PARTN
```

This will create a special file called PARTNSAV.FIL on a floppy disk you designate. This file contains an image of all of the boot sectors across all of the DOS accessible partitions on your hard disk. This information can later be restored by the UNFORMAT command as follows:

```
UNFORMAT /PARTN
```

The UNFORMAT program will ask for the disk containing the PARTNSAV.FIL file and restore the boot sector information to the hard disk.

MIRROR can also be used to backup the FATs and Directory structure to a special hidden file on the disk by simply executing the command with no parameters at all:

```
MIRROR
```

By executing the command in this manner, a special file called MIRROR.FIL, as well as a hidden file called MIRRORSAV.FIL, will be created in the root directory, and a copy of the FAT and Directory structures will be copied into some free space at the end of the drive. The actual FAT and Directory data is not stored in a normal file with a filename but is instead written to empty sectors at the end of the drive. These sectors will eventually be overwritten, which is why it is a good idea to run the MIRROR command frequently. The second time you run it, it will actually create a second backup and retain the first. Every subsequent run will retain the previous one as a secondary backup. It is a good idea to put the MIRROR command in your AUTOEXEC.BAT file so that these critical areas of the disk are backed up every time your system boots. Usually, MIRROR only takes a few seconds to perform its backup functions so there will be little delay in booting.

To restore the FAT and Directory area backups, you would use the UNFORMAT program, also with no parameters. This program will then search the disk for the MIRROR backups and prompt you for a possible restore. Be careful in restoring these areas as they may be out of sync somewhat with the actual files on your disk. You should always run MIRROR after defragmenting your drive because the defragmenting process moves many files on the disk, making the MIRROR backup obsolete.

There is one unfortunate problem with the MIRROR program; both Microsoft and IBM removed it from DOS 6. Fortunately, you can retain your copy of MIRROR from DOS 5, or you can download the DOS 6 supplemental disk from the MSDOS forum on CompuServe. This package will provide you with a number of utilities left out of DOS 6 but the MIRROR command is the most useful. You can also order this disk directly from Microsoft for $5 using an order form in the back of the DOS 6 manual from Microsoft.

Many other data recovery programs offer the ability to back up these areas on the disk. The Norton Utilities by Symantec is one of the best and most well known data recovery packages on the market. It offers this capability and does a much better job than DOS alone. Norton uses an IMAGE command instead of MIRROR, and also has a special program called RESCUE that creates a rescue disk with not only the boot sector areas on it, but copies of your AUTOEXEC.BAT and CONFIG.SYS files, the DOS system files, and the CMOS RAM data from your system. Also included on the rescue disk are copies of the appropriate Norton Utilities programs that can be used to restore these areas.

Checking for Virus Programs. Both Microsoft and IBM now provide standard anti-virus software in DOS 6.x. The Microsoft Anti-Virus program is actually a reduced function version of the Central Point Anti-Virus software. IBM has written a package called the IBM Anti-Virus program. Many aftermarket utility packages are available that will scan for and remove virus programs. One of the best known is the McAfee Associates SCANV program, which is also one of the easiest to run because it is a command line utility. The McAfee program is also distributed through BBS systems and is often site-licensed to large companies.

No matter which of these programs you use, it is a good idea to perform a scan for virus programs periodically, especially before making hard disk backups. This will help to ensure that you catch any potential virus problem before it spreads and becomes a major hassle.

Reformatting a Hard Disk. Periodically reformatting a hard disk is an operation often overlooked as part of a preventive maintenance plan. Reformatting serves two purposes. On non-servo controlled drives (stepper motor head actuators) the low level format rewrites the sector header information in alignment with current head positions, which can drift in stepper motor head actuator drives because of temperature- and stress-induced dimensional changes between the platters and heads. If these alignment variations continue unchecked, they eventually will cause read and write errors. Note that this reformatting operation applies to stepper motor head actuator drives only and not to voice-coil drives, which maintain their positional accuracy due to the closed loop servo control mechanism that guides the heads.

The second function of a Low Level Format is to locate and mark or spare out any new defective sectors. This can be accomplished during the Low Level Format and also by a subsequent surface analysis.

Reformatting a hard disk lays down new track and sector ID marks and boundaries, remarks the manufacturer's defects, and performs a surface scan for new defects that might have developed since the last format. With Zone Recorded drives, only the defect mapping and surface analysis are performed. The sector headers are usually not completely rewritten. Temperature variations, case flexing, and physical positioning can add up to eventual read and write errors in a stepper motor head actuated hard disk. This type of failure sometimes appears as a gradually increasing number of disk retries and other read and write problems. You also might notice difficulties when you boot the disk for the first time each day or if the system has been turned off for some time. The cause of these problems is a mistracking between where the data has actually written on the drive and where the track and sector ID marks are located. If the drive is used in a variety of temperatures and environmental conditions, dimensional changes between the heads and platters can cause the data to be written at various improper offsets from the desired track locations.

The reformatting procedure for hard disks is the equivalent of aligning a floppy drive. For hard disks, however, the concern is not for the actual location of each track, but that the drive heads are positioned accurately to the same track location each time. Because of the inherent problems in tracking with stepper motor drives and the lack of a track-following system, mistracking errors accumulate and eventually cause a failure to read or write a particular location. To correct this problem, you must lay down a new set of track and sector ID marks that correspond as closely as possible to the position from which the heads actually read and write data. To do this (with stepper motor drives), you must perform a low-level format.

To make the new format effective, you must do it at the drive's full operating temperature and with the drive in its final mounted position. If the drive runs on its side when it is installed, the format must be done in that position.

For inexpensive drives that lack a proper shock-mounting system (such as the Seagate ST-225, -238, -251, or any ST-2XX drive), make sure that the drive is completely installed before beginning low-level formatting. When you attach the mounting screws to these drives, you are placing screws almost directly into the Head Disk Assembly (HDA), which can cause the HDA to bend or warp slightly depending on how much you tighten the screws. Turn the screws just until they are snug. Do not over tighten them, or your drive might fail if the screws ever loosen and the HDA stress is relaxed. Screws that are too tight can cause continuous read and write problems from stress in the HDA. Having this type of drive completely installed when you are formatting places the HDA under the same physical stress and distortion that it will be under when data is read and written, which makes the format much more accurate.

The frequency with which you should reformat a hard disk depends primarily on the types of drives you have. If the drives are inexpensive stepper motor types (the Seagate ST-2XX series, for example) you probably should reformat the drives once a year. People who must support large numbers of these cheaper drives become known as "hard disk reformatting specialists." A joke in the industry is that some of these drives require winter and summer formats because of temperature sensitivity. This joke, unfortunately, can be somewhat truthful in some cases.

High-quality voice-coil drives usually are formatted only once, either at the factory, as in the case of most IDE or SCSI drives, or by the installer, as is the case with most other types of hard disks. With these drives, a reformat is only performed when the drive begins to exhibit problems reading or writing any sectors on the disk. This might appear in the form of DOS “Abort, Retry, Fail, Ignore” error messages, or other read or write errors. Upon encountering difficulty with any sectors on the disk, it should be fully backed up and then a reformat should be performed. In this case, the reformat and subsequent surface analysis will locate and mark off or spare out the marginal sectors, thus restoring the drive to proper operation.

As mentioned earlier, voice-coil drives do not require reformatting the hard disk as do stepper motor drives. Voice-coil drives do not usually develop difficulties with *hysteresis,* a measurement of how accurately a drive can repeatedly locate to a specified position. Hysteresis is measured by commanding a drive to position itself at a particular cylinder and later (at a different temperature), commanding the drive to position itself at the same cylinder. The voice-coil drive always positions itself at the same position relative to the disk platter because of the track-following servo guide head. The stepper motor drive, however, is fooled by temperature and other environmental or physical stress changes because it is essentially a *blind* positioning system.

Refer to Chapter 14, “Hard Disk Drives and Controllers,” which describes hard disk formatting procedures, for more information about the proper tools and procedures for reformatting the different types of hard disk drives.

Passive Preventive Maintenance Procedures

Passive preventive maintenance involves taking care of the system in an external manner: basically, providing the best possible environment—both physical as well as electrical—for the system to operate in. Physical concerns are conditions such as ambient temperature, thermal stress from power cycling, dust and smoke contamination, and disturbances such as shock and vibration. Electrical concerns are items such as electrostatic discharge (ESD), power line noise, and radio-frequency interference. Each of these environmental concerns is discussed in this section.

Examining the Operating Environment. Oddly enough, one of the most overlooked aspects of microcomputer preventive maintenance is protecting the hardware—and the sizable financial investment it represents—from environmental abuse. Computers are relatively forgiving, and they generally are safe in an environment that is comfortable for people. Computers, however, often are treated with no more respect than desktop calculators. The result of this type of abuse is a system failure.

Before you acquire a system, prepare a proper location for your new system, free of airborne contaminants such as smoke or other pollution. Do not place your system in front of a window: the system should not be exposed to direct sunlight or temperature variations. The environmental temperature should be as constant as possible. Power should be provided through properly grounded outlets, and should be stable and free from electrical noise and interference. Keep your system away from radio transmitters or other sources of radio frequency energy. This section examines these issues in more detail.

Heating and Cooling. Thermal expansion and contraction from temperature changes place stress on a computer system. Therefore, keeping the temperature in your office or room relatively constant is important to the successful operation of your computer system.

Temperature variations can lead to serious problems. You might encounter excessive chip creep, for example. If extreme variations occur over a short period, signal traces on circuit boards can crack and separate, solder joints can break, and contacts in the system can undergo accelerated corrosion. Solid-state components such as chips can be damaged also, and a host of other problems can develop.

Temperature variations can play havoc with hard disk drives also. Writing to a disk at different ambient temperatures can, on some drives, cause data to be written at different locations relative to the track centers. Read and write problems then might accelerate later.

To ensure that your system operates in the correct ambient temperature, you first must determine your system's specified functional range. Most manufacturers provide data about the correct operating temperature range for their systems. Two temperature specifications might be available, one indicating allowable temperatures during operation and another indicating allowable temperatures under non-operating conditions. IBM, for example, indicates the following temperature ranges as acceptable for most of its systems:

> System on: 60 to 90 degrees Fahrenheit
>
> System off: 50 to 110 degrees Fahrenheit

For the safety of the disk and the data it contains, avoid rapid changes in ambient temperatures. If rapid temperature changes occur—for example, when a new drive is shipped to a location during the winter and then brought indoors—let the drive acclimate to room temperature before turning it on. In extreme cases, condensation forms on the platters inside the drive head disk assembly—disastrous for the drive if you turn it on before the condensation can evaporate. Most drive manufacturers specify a timetable to use as a guide in acclimating a drive to room temperature before operating it. You usually must wait several hours to a day before a drive is ready to use after it has been shipped or stored in a cold environment.

Most office environments provide a stable temperature in which to operate a computer system, but some do not. Be sure to give some consideration to the placement of your equipment.

Power Cycling (On/Off). As you have just learned, the temperature variations a system encounters greatly stress the system's physical components. The largest temperature variations a system encounters, however, are those that occur during system warm-up when you initially turn it on. Turning on (also called powering on) a cold system subjects it to the greatest possible internal temperature variations. For these reasons, limiting the number of power-on cycles a system is exposed to greatly improves its life and reliability.

If you want a system to have the longest, most trouble-free life possible, you should limit the temperature variations in its environment. You can limit the extreme temperature cycling in two simple ways during a cold start-up: leave the system off all the time or leave it on all the time. Of these two possibilities, of course, you want to choose the latter option. Leaving the power on is the best way I know to promote system reliability. If your only concern is system longevity, the simple recommendation would be to keep the system unit powered on (or off!) continuously. In the real world, however there are more variables to consider, such as the cost of electricity, the potential fire hazard of unattended running equipment, and other concerns as well.

If you think about the way light bulbs typically fail, you can begin to understand that thermal cycling can be dangerous. Light bulbs burn out most often when you first turn them on, because the filament must endure incredible thermal stress as it changes temperature—in less than one second—from ambient to several thousands of degrees. A bulb that remains on continuously lasts longer than one that is turned on and off repeatedly.

Some people argue that the reason you should leave a computer system on continuously is to prevent the electrical "shock" from the inrush of power when you start up a system. The cause of failure in a low-voltage solid-state circuit repeatedly powered on and off, however, is not in-rushing electrons, but rather physical stresses caused by thermal expansion and contraction of the components. Component engineers agree, and tests prove, that a device left on continuously outlasts one that is powered on and off repeatedly.

Where problems can occur immediately at power-on is in the power supply. The start-up current draw for the system and for any motor during the first few seconds of operation is very high compared to the normal operating-current draw. Because the current must come from the power supply, the supply has an extremely demanding load to carry for the first few seconds of operation, especially if several disk drives will be started. Motors have an extremely high power-on current draw. This demand often overloads a marginal circuit or component in the supply and causes it to burn or break with a "snap." I have seen several power supplies die the instant a system was powered up. To enable your equipment to have the longest possible life, try to keep the temperature of solid-state components relatively constant, and limit the number of start-ups on the power supply. The only way I know to do so is to leave the system on.

Although it sounds as though I am recommending that you leave all of your computer equipment on 24 hours a day, seven days a week, I no longer recommend this type of operation. A couple of concerns have tempered my urge to leave everything running continuously. One is that an unattended system that is powered on, represents a fire hazard. I have seen monitors start themselves on fire after internally shorting, and systems whose cooling fans have frozen, enabling the power supply and entire system to overheat. I do not leave any system running in an unattended building. Another problem is wasted electrical power. Many companies have adopted austerity programs that involve turning lights and other items off when not in use. The power consumption of some of today's high-powered systems and accessories is not trivial. Also, an unattended operating system is more of a security risk than one that is powered off and locked.

Realities—such as the fire hazard of unattended systems running during night or weekend hours, security problems, and power-consumption issues—might prevent you from leaving your system on all the time. Therefore, you must compromise. Power on the system only one time daily. Don't power the system on and off several times every day. This good advice is often ignored, especially when several users share systems. Each user powers on the system to perform work on the PC and then powers off the system. These systems tend to have many more problems with component failures.

If you are concerned about running your hard disk continuously, let me dispel your fears. Running your hard disk continuously might be the best thing you can do for your drive. Leaving the drive powered on is the best method for reducing read and write failures caused by temperature changes. If you are using extremely inexpensive drives with stepper motor actuators, leaving the drive on greatly improves reliability and increases the time between low-level formats caused by mistracking. A drive's bearings and motors also function longer if you reduce the power-on temperature cycling. You might have had a disk that didn't boot after you turned the drive off for a prolonged period (over the weekend, for example) and you fixed the problem with a subsequent low-level format, but you most likely wouldn't have had the problem if you had left your drive on.

If you are in a building with a programmable thermostat, you have another reason to be concerned about temperatures and disk drives. Some buildings have thermostats programmed to turn off the heat overnight or over the weekend. These thermostats are programmed also to quickly raise the temperature just before business hours every day. In Chicago, for example, outside temperatures in the winter can dip to 20 degrees below 0 (not including a wind-chill factor). An office building's interior temperature can drop as low as 50 degrees during the weekend. When you arrive Monday morning, the heat has been on for only an hour or so, but the hard disk platters might have not yet reached even 60 degrees when you turn on the system unit. During the first 20 minutes of operation, the disk platters rapidly rise in temperature to 120 degrees or more. If you have an inexpensive stepper motor hard disk and begin writing to the disk at these low temperatures, you are setting yourself up for trouble. Again, many systems with these "cheap" drives don't even boot properly under these circumstances and must be warmed up before they even boot DOS.

Tip

If you do not leave a system on continuously, at least give it 15 minutes or more to warm up before writing to the hard disk. Power up the system and go get a cup of coffee, read the paper, or do some other task. This practice does wonders for the reliability of the data on your disk, especially cheaper units.

If you do leave your system on for long periods of time, make sure that the screen is blank or displays a random image if the system is not in use. The phosphor on the picture tube can burn if a stationary image is left on-screen continuously. Higher-persistence phosphor monochrome screens are most susceptible, and the color displays with low-persistence phosphors are the least susceptible. If you ever have seen a monochrome display with the image of some program permanently burned in—even with the display off—you know what I mean. Look at the monitors that display flight information at the airport—they usually show the effects of phosphor burn.

Screen savers or blankers will either blank the screen completely or display some sort of moving random image to prevent burn in. This can be accomplished by either a manual or automatic procedure as follows:

- *Manual.* Turn the brightness and contrast levels all the way down, or even power the display off completely. This technique is effective but it is a manual method; you must remember to do it.
- *Automatic.* Many types of programs can cause the screen to blank or display random images automatically at a predetermined interval. Screen savers are built into most Graphical User Interfaces (GUIs) such as Windows and OS/2. These can easily be enabled, and you can also specify the time delay before they activate. If you run under plain DOS, you can use a number of public domain as well as commercial screen saver programs. These programs usually run as terminate-and-stay-resident (TSR) programs. The program watches the clock as well as the keyboard and mouse ports. If several minutes pass with nothing typed at the keyboard or no mouse movement, the program activates and either shuts off all signals to the display or creates an image that moves around on the screen to prevent burn in.

Static Electricity. Static electricity can cause numerous problems within a system. The problems usually appear during the winter months when humidity is low or in extremely dry climates where the humidity is low year-round. In these cases, you might need to take special precautions to ensure that the system functions properly.

Static discharges outside a system-unit chassis are rarely a source of permanent problems within the system. The usual effect of a static discharge to the case, keyboard, or even in close proximity to a system, is a parity check (memory) error or a locked-up system. In some cases, I have been able to cause parity checks or system lockups by simply walking past a system. Most static-sensitivity problems such as this one are caused by improper grounding of the system power. Be sure that you always use a three-prong, grounded power cord plugged into a properly grounded outlet. If you are unsure about the outlet, you can buy an outlet tester at most electronics-supply or hardware stores for only a few dollars.

Whenever you open a system unit or handle circuits removed from the system, you must be much more careful with static. You can damage permanently a component with a static discharge if the charge is not routed to a ground. I usually recommend handling boards and adapters first by a grounding point such as the bracket to minimize the potential for static damage.

An easy way to prevent static problems is with good power-line grounding, which is extremely important for computer equipment. A poorly designed power-line grounding system is one of the primary causes of poor computer design. The best way to prevent static damage is to prevent the static charge from getting into the computer in the first place. The chassis ground in a properly designed system serves as a static guard for the computer, which redirects the static charge safely to ground. For this ground to be complete, therefore, the system must be plugged into a properly grounded three-wire outlet.

If the static problem is extreme, you can resort to other measures. One is to use a grounded static mat underneath the computer. Touch the mat first before you touch the computer, to ensure that any static charges are routed to ground and away from the system unit's internal parts. If problems still persist, you might want to check out the electrical building ground. I have seen installations in which three-wire outlets exist but are not grounded properly. You can use an outlet tester to be sure that the outlet is wired properly.

Power Line Noise. To run properly, a computer system requires a steady supply of clean, noise-free power. In some installations, however, the power line serving the computer serves heavy equipment also, and the voltage variations resulting from the on-off cycling of this equipment can cause problems for the computer. Certain types of equipment on the same power line also can cause voltage *spikes*—short transient signals of sometimes 1,000 volts or more—that can physically damage a computer. Although these spikes are rare, they can be crippling. Even a dedicated electrical circuit used only by a single computer can experience spikes and transients, depending on the quality of the power supplied to the building or circuit.

During the site-preparation phase of a system installation, you should be aware of these factors to ensure a steady supply of clean power:

- If possible, the computer system should be on its own circuit with its own circuit breaker. This setup does not guarantee freedom from interference but it helps.
- The circuit should be checked for a good, low-resistance ground, proper line voltage, freedom from interference, and freedom from brownouts (voltage dips).
- A three-wire circuit is a must but some people substitute grounding-plug adapters to adapt a grounded plug to a two-wire socket. This setup is not recommended; the ground is there for a reason.
- Power line noise problems increase with the resistance of the circuit, which is a function of wire size and length. To decrease resistance, therefore, avoid extension cords unless absolutely necessary, and then use only heavy-duty extension cords.
- Inevitably, you will want to plug in other equipment later. Plan ahead to avoid temptations to use too many items on a single outlet. If possible, provide a separate power circuit for noncomputer-related accessories.

Air conditioners, coffee makers, copy machines, laser printers, space heaters, vacuum cleaners, and power tools are some of the worst corrupters of a PC system's power. Any of these items can draw an excessive amount of current and play havoc with a PC system on the same electrical circuit. I've seen offices in which all the computers begin to crash at about 9:05 a.m. daily, which is when all the coffee makers are turned on!

Also, try to ensure that copy machines and laser printers do not share a circuit with other computer equipment. These devices draw a large amount of power.

Another major problem in some companies is partitioned offices. Many of these partitions are prewired with their own electrical outlets and are plugged into one another in a sort of power line daisy chain, similar to chaining power strips together. I pity the person in the cubicle at the end of the electrical daisy chain, who will have very flaky power!

As a real-world example of too many devices sharing a single circuit, I can describe several instances in which a personal computer had a repeating parity check problem. All efforts to repair the system had been unsuccessful. The reported error locations from the parity check message also were inconsistent, which normally indicates a problem with power. The problem could have been the power supply in the system unit or the external power supplied from the wall outlet. This problem was solved one day as I stood watching the system. The parity check message was displayed at the same instant someone two cubicles away turned on a copy machine. Placing the computers on a separate line solved the problem.

By following the guidelines in this section, you can create the proper power environment for your systems and help to ensure trouble-free operation.

Radio-Frequency Interference. Radio-frequency interference (RFI) is easily overlooked as a problem factor. The interference is caused by any source of radio transmissions near a computer system. Living next door to a 50,000-watt commercial radio station is one sure way to get RFI problems, but less powerful transmitters cause problems too. I know of many instances in which portable radio-telephones have caused sporadic random keystrokes to appear, as though an invisible entity were typing on the keyboard. I also have seen RFI cause a system to lock up. Solutions to RFI problems are more difficult to state because every case must be handled differently. Sometimes, reorienting a system unit eliminates the problem because radio signals can be directional in nature. At other times, you must invest in specially shielded cables for cables outside the system unit, such as the keyboard cable.

One type of solution to an RFI noise problem with cables is to pass the cable through a *toroidal iron core,* a doughnut-shaped piece of iron placed around a cable to suppress both the reception and transmission of electromagnetic interference (EMI). If you can isolate an RFI noise problem in a particular cable, you often can solve the problem by passing the cable through a toroidal core. Because the cable must pass through the center hole of the core, it often is difficult, if not impossible, to add a toroid to a cable that already has end connectors installed.

Radio Shack sells a special snap-together toroid designed specifically to be added to cables already in use. This toroid looks like a thick-walled tube that has been sliced in half. You just lay the cable in the center of one of the halves, and snap the other half over the first. This type of construction makes it easy to add the noise-suppression features of a toroid to virtually any existing cable.

IBM also makes a special 6-foot long PS/2 keyboard cable with a built-in toroid core (part number 27F4984) that can greatly reduce interference problems. This cable has the smaller 6-pin DIN (PS/2 style) connector at the system end and the standard SDL (Shielded Data Link) connector at the keyboard end; it costs about $40.

The best, if not the easiest, way to eliminate the problem probably is to correct it at the source. You likely won't convince the commercial radio station near your office to shut down but if you are dealing with a small radio transmitter that is generating RFI, sometimes you can add to the transmitter a filter that suppresses spurious emissions. Unfortunately, problems sometimes persist until the transmitter is either switched off or moved some distance away from the affected computer.

Note that your own computer systems can be a source of RFI. Computer equipment must meet one of these two classifications to be certified and salable, according to the FCC (Federal Communications Commission): Class A or Class B. The Class A specification applies to computing devices sold for use in commercial, business, and industrial environments. Class B indicates that the equipment has passed more stringent tests and can be used in residential environments, in addition to any environments allowed under Class A.

The FCC does not really police users or purchasers of computer equipment so much as it polices the equipment manufacturers or vendors. Therefore, if you are using a Class A-rated system in your home, you don't need to worry about radio police showing up at your door. You must obtain a Class B certification for systems that meet one of the following conditions:

- Marketed through retail or direct-mail outlets
- Sold to the general public rather than commercial users only
- Operates on battery or 120-volt AC electrical power

Notice that a system has to fit each of these three categories to be considered a personal computer and be subject to the stricter Class B rules. Notice also that the FCC considers all portable computer systems as meeting Class B standards because their portability makes them likely to be used in a residential setting.

The FCC standards for Class A and Class B certification govern two kinds of emissions: conductive emissions, radiated from the computer system into the power cord, and radio-frequency emissions, radiated from the computer system into space. Table 18.1 shows the conductive and radio-frequency emissions limitations to be eligible for both Class A and Class B ratings.

Table 18.1 FCC Class A and Class B Emission Limitations

Conductive Emissions: Maximum Signal Level (mv)

Frequency	Class A	Class B
0.45 to 1.705 MHz	1000	250
1.705 to 30.0 MHz	3000	250

Radiated Emissions: Maximum Field Strength (uv/M)

Frequency	Class A (10M)	Class B (3M0
30 to 88 MHz	90	100
88 to 216 MHz	150	150
217 to 960 MHz	210	200
960 MHz and up	300	500

MHz = Megahertz
10M = Measured at 10 meters
3M = Measured at 3 meters
mv = Millivolts
uv/M = Microvolts per meter

Notice that although some of the specific numbers listed for Class A seem lower than those required for Class B, you must consider that field strength measurements normally decline under the inverse square law: the strength of the signal decreases as the square of the distance from the source. A rating of 100 microvolts per meter at 3 meters, therefore, would be approximately equal to a rating of about 9 microvolts per meter at 10 meters. This calculation just means that the limits to pass Class B certification are much tougher than they look, and certainly are much tougher than Class A limits. Additionally, Class A certification is tested and verified entirely by the manufacturer; Class B certification requires a sample of the equipment to be sent to the FCC for testing.

IBM and most other responsible manufacturers ensure that all systems they sell meet the stricter Class B designations. One of the primary reasons for the Micro Channel Architecture design was to meet, and also greatly exceed these FCC classifications. IBM knew that as computing clock speeds go up, so do the radio emissions. As clock rates of 66, 75, and 100 MHz and higher become more common (and they will soon), IBM will have a distinct advantage over manufacturers still using the AT or EISA bus designs because vendors using ISA or EISA bus designs will have to invest in more expensive chassis and case shielding to combat the emission problem. IBM and other Micro Channel clones will gain a distinct manufacturing cost advantage.

Dust and Pollutants. Dirt, smoke, dust, and other pollutants are bad for your system. The power-supply fan carries airborne particles through your system, and they collect inside. If your system is used in an extremely harsh environment, you might want to investigate some of the industrial systems on the market designed for harsh conditions.

IBM used to sell industrial model XT and AT systems but discontinued them after introducing the PS/2. IBM has licensed several third-party companies to produce industrial versions of PS/2 systems.

Compatible vendors also have industrial systems; many companies make special *hardened* versions of their systems for harsh environments. Industrial systems usually use a different cooling system from the one used in a regular PC. A large cooling fan is used to pressurize the case rather than depressurize it, as most systems do. The air pumped into the case passes through a filter unit that must be cleaned and changed periodically. The system is pressurized so that no contaminated air can flow into it; air flows only outward. The only way air can enter is through the fan and filter system.

These systems also might have special keyboards impervious to liquids and dirt. Some flat-membrane keyboards are difficult to type on, but are extremely rugged; others resemble the standard types of keyboards, but have a thin, plastic membrane that covers all the keys. You can add this membrane to normal types of keyboards to seal them from the environment.

A new breed of humidifier can cause problems with computer equipment. This type of humidifier uses ultrasonics to generate a mist of water sprayed into the air. The extra humidity helps cure problems with static electricity resulting from a dry climate, but the airborne water contaminants can cause many problems. If you use one of these systems, you might notice a white ash-like deposit forming on components. The deposit is the result of abrasive and corrosive minerals suspended in the vaporized water. If these deposits collect on the disk drive heads, they will ruin the heads and scratch disks. The only safe way to run one of these ultrasonic humidifiers is with pure distilled water. If you use a humidifier, be sure it does not generate these deposits.

If you do your best to keep the environment for your computer equipment clean, your system will run better and last longer. Also, you will not have to open your unit as often for complete preventive maintenance cleaning.

Using Power-Protection Systems

Power-protection systems do just what the name implies: they protect your equipment from the effects of power surges and power failures. In particular, power surges and spikes can damage computer equipment, and a loss of power can result in lost data. In this section, you learn about the four primary types of power-protection devices available and under what circumstances you should use them.

Before considering any further levels of power protection, you should know that the power supply in your system (if your system is well-made) already affords you a substantial amount of protection. The power supplies in IBM equipment are designed to provide protection from higher-than-normal voltages and currents, and provide a limited amount of power line noise filtering. Some of the inexpensive aftermarket power supplies probably do not have this sort of protection; be careful if you have an inexpensive clone system. In those cases, further protecting your system might be wise.

IBM's PS/2 power supplies will stay within operating specifications and continue to run a system if any of these power line disturbances occur:

- Voltage drop to 80 volts for up to 2 seconds
- Voltage drop to 70 volts for up to .5 seconds
- Voltage surge of up to 143 volts for up to 1 second

IBM also states that neither their power supplies nor systems will be damaged by the following occurrences:

- Full power outage
- Any voltage drop (brownout)
- A spike of up to 2,500 volts

Because of the high quality power supply design that IBM uses, they state in their documentation that external surge suppressors are not needed for PS/2 systems. Most other high quality name brand manufacturers also use high quality power supply designs. Companies like Astec, PC Power and Cooling, and others make very high quality units.

To verify the levels of protection built into the existing power supply in a computer system, an independent laboratory subjected several unprotected PC systems to various spikes and surges up to *6,000 volts*—considered the maximum level of surge that can be transmitted to a system by an electrical outlet. Any higher voltage would cause the power to arc to ground within the outlet itself. Note that none of the systems sustained permanent damage in these tests; the worst thing that happened was that some of the systems rebooted or shut down if the surge was more than 2,000 volts. Each system restarted when the power switch was toggled after a shutdown.

I do not use any real form of power protection on my systems, and they have survived near-direct lightning strikes and powerful surges. The most recent incident, only 50 feet from my office, was a direct lightning strike to a brick chimney that *blew the top of the chimney apart*. None of my systems (which were running at the time) were damaged in any way from this incident; they just shut themselves down. I was able to restart each system by toggling the power switches. An alarm system located in the same office, however, was destroyed by this strike. I am not saying that lightning strikes or even much milder spikes and surges cannot damage computer systems—another nearby lightning strike did destroy a modem and serial adapter installed in one of my systems. I was just lucky that the destruction did not include the motherboard.

This discussion points out an important oversight in some power-protection strategies: you may elect to protect your systems from electrical power disturbances, but do not forget to provide similar protection also from spikes and surges on the phone line.

The automatic shutdown of a computer during power disturbances is a built-in function of most high-quality power supplies. You can reset the power supply by flipping the power switch from on to off and back on again. Some power supplies, such as those in most of the PS/2 systems, have an *auto-restart function*. This type of power supply acts the same as others in a massive surge or spike situation: it shuts down the system. The difference is that after normal power resumes, the power supply waits for a specified delay of three to six seconds and then resets itself and powers the system back up. Because no manual switch resetting is required, this feature is desirable in systems functioning as a network file server or in a system in a remote location.

The first time I witnessed a large surge cause an immediate shutdown of all my systems, I was extremely surprised. All the systems were silent, but the monitor and modem lights were still on. My first thought was that everything was blown, but a simple toggle of each system-unit power switch caused the power supplies to reset, and the units powered up with no problem. Since that first time, this type of shutdown has happened to me several times, always without further problems.

The following types of power-protection devices are explained in the sections that follow:

- Surge suppressors
- Line conditioners
- Standby power supplies (SPS)
- Uninterruptible power supplies (UPS)

Surge Suppressors (Protectors)

The simplest form of power protection is any of the commercially available surge protectors; that is, devices inserted between the system and the power line. These devices, which cost between $20 and $200, can absorb the high-voltage transients produced by nearby lightning strikes and power equipment. Some surge protectors can be effective for certain types of power problems, but they offer only very limited protection.

Surge protectors use several devices, usually metal-oxide varistors (MOVs), that can clamp and shunt away all voltages above a certain level. MOVs are designed to accept voltages as high as 6,000 volts and divert any power above 200 volts to ground. MOVs can handle normal surges, but powerful surges such as a direct lightning strike can blow right through them. MOVs are not designed to handle a very high level of power, and self-destruct while shunting a large surge. These devices, therefore, cease to function after a single large surge or a series of smaller ones. The real problem is that you cannot easily tell when they no longer are functional; the only way to test them is to subject the MOVs to a surge, which destroys them. Therefore, you never really know if your so-called surge protector is protecting your system.

Some surge protectors have status lights that let you know when a surge large enough to blow the MOVs has occurred. A surge suppressor without this status indicator light is useless because you never know when it has stopped protecting.

Underwriters Laboratories has produced an excellent standard that governs surge suppressors, called UL 1449. Any surge suppressor that meets this standard is a very good one, and definitely offers an additional line of protection beyond what the power supply in your PC already does. The only types of surge suppressors worth buying, therefore, should have two features: conformance to the UL 1449 standard and a status light indicating when the MOVs are blown. Units that meet the UL 1449 specification say so on the packaging or directly on the unit. If this standard is not mentioned, it does not conform, and you should avoid it.

Another good feature to have in a surge suppressor is a built-in circuit breaker that can be reset rather than a fuse. The breaker protects your system if the system or a peripheral develops a short. These better surge suppressors usually cost about $40.

Phone Line Surge Protectors

In addition to protecting the power lines, it is critical to provide protection to your systems from any phone lines that are connected. If you are using a modem or fax board which is plugged into the phone system, any surges or spikes that travel the phone line can potentially damage your system. In many areas, the phone lines are especially susceptible to lightning strikes, which is the largest cause of fried modems and computer equipment attached to them.

Several companies manufacture or sell simple surge protectors that plug between your modem and the phone line. These inexpensive devices can be purchased from most electronics supply houses. Most of the cable and communication product vendors listed in Appendix B sell these phone line surge protectors.

Line Conditioners

In addition to high-voltage and current conditions, other problems can occur with incoming power. The voltage might dip below the level needed to run the system and result in a brownout. Other forms of electrical noise other than simple voltage surges or spikes might be on the power line, such as radio-frequency interference or electrical noise caused by motors or other inductive loads.

Remember two things when you wire together digital devices (such as computers and their peripherals). A wire is an antenna and has a voltage induced in it by nearby electromagnetic fields, which can come from other wires, telephones, CRTs, motors, fluorescent fixtures, static discharge, and, of course, radio transmitters. Digital circuitry also responds with surprising efficiency to noise of even a volt or two, making those induced voltages particularly troublesome. The wiring in your building can act as an antenna and pick up all kinds of noise and disturbances. A line conditioner can handle many of these types of problems.

A line conditioner is designed to remedy a variety of problems. It filters the power, bridges brownouts, suppresses high-voltage and current conditions, and generally acts as a buffer between the power line and the system. A line conditioner does the job of a surge suppressor, and much more. It is more of an active device functioning continuously rather than a passive device that activates only when a surge is present. A line conditioner provides true power conditioning and can handle myriad problems. It contains transformers, capacitors, and other circuitry that temporarily can bridge a brownout or low-voltage situation. These units usually cost several hundreds of dollars, depending on the power-handling capacity of the unit.

Backup Power

The next level of power protection includes backup power-protection devices. These units can provide power in case of a complete blackout, which provides the time needed for an orderly system shutdown. Two types are available: the standby power supply (SPS) and the uninterruptible power supply (UPS). The UPS is a special device because it does much more than just provide backup power: it is also the best kind of line conditioner you can buy.

Standby Power Supplies (SPS). A standby power supply is known as an *off-line device*: it functions only when normal power is disrupted. An SPS system uses a special circuit that can sense the AC line current. If the sensor detects a loss of power on the line, the system quickly switches over to a standby battery and power inverter. The power inverter converts the battery power to 110-volt AC power, which then is supplied to the system.

SPS systems do work, but sometimes a problem occurs with the switch to battery power. If the switch is not fast enough, the computer system unit shuts down or reboots anyway, which defeats the purpose of having the backup power supply. A truly outstanding SPS adds to the circuit a *ferroresonant transformer,* a large transformer with the capability to store a small amount of power and deliver it during the switch time. Having this device is similar to having on the power line a buffer that you add to an SPS to give it almost truly uninterruptible capability.

SPS units also may or may not have internal line conditioning of their own; most cheaper units place your system directly on the regular power line under normal circumstances and offer no conditioning. The addition of a ferroresonant transformer to an SPS gives it additional regulation and protection capabilities due to the buffer effect of the transformer. SPS devices without the ferroresonant transformer still require the use of a line conditioner for full protection. SPS systems usually cost from $200 to several thousands of dollars, depending on the quality and power-output capacity.

Uninterruptible Power Supplies (UPS). Perhaps the best overall solution to any power problem is to provide a power source that is both conditioned and that also cannot be interrupted—which describes an uninterruptible power supply. UPSs are known as online systems because they continuously function and supply power to your computer systems. Because some companies advertise ferroresonant SPS devices as though they were UPS devices, many now use the term *true UPS* to describe a truly on-line system. A true UPS system is constructed much the same as an SPS system; however, because you always are operating from the battery, there is no switching circuit.

In a true UPS, your system always operates from the battery, with a voltage inverter to convert from 12 volts DC to 110 volts AC. You essentially have your own private power system that generates power independently of the AC line. A battery charger connected to the line or wall current keeps the battery charged at a rate equal to or greater than the rate at which power is consumed.

When power is disconnected, the true UPS continues functioning undisturbed because the battery-charging function is all that is lost. Because you already were running off the battery, no switch takes place and no power disruption is possible. The battery then begins discharging at a rate dictated by the amount of load your system places on the unit, which (based on the size of the battery) gives you plenty of time to execute an orderly system shutdown. Based on an appropriately scaled storage battery, the UPS functions continuously, generating power and preventing unpleasant surprises. When the line power returns, the battery charger begins recharging the battery, again with no interruption.

UPS cost is a direct function of both the length of time it can continue to provide power after a line current failure, and how much power it can provide; therefore, purchasing a UPS that gives you enough power to run your system and peripherals as well as enough time to close files and provide an orderly shutdown would be sufficient. In most PC applications, this solution is the most cost-effective because the batteries and charger portion of the system must be much larger than the SPS type of device, and will be more costly.

Many SPS systems are advertised as though they were true UPS systems. The giveaway is the unit's "switch time." If a specification for switch time exists, the unit cannot be a true UPS because UPS units never switch. Understand, however, that a good SPS with a ferroresonant transformer can virtually equal the performance of a true UPS at a lower cost.

Because of a UPS's almost total isolation from the line current, it is unmatched as a line conditioner and surge suppressor. The best UPS systems add a ferroresonant transformer for even greater power conditioning and protection capability. This type of UPS is the best form of power protection available. The price, however, can be very high. A true UPS costs from $1 to $2 per watt of power supplied. To find out just how much power your system requires, look at the UL sticker on the back of the unit. This sticker lists the maximum power draw in watts, or sometimes in just volts and amperes. If only voltage and amperage are listed, multiply the two figures to calculate a wattage figure.

As an example, the back of an IBM PC AT Model 339 indicates that the system can require as much as 110 volts at a maximum current draw of 5 amps. The maximum power this AT can draw is about 550 watts. This wattage is for a system with every slot full, two hard disks and one floppy—in other words, the maximum possible level of expansion. The system should never draw any more power than that; if it does, a 5-ampere fuse in the power supply blows. This type of system normally draws an average 300 watts; to be safe when you make calculations for UPS capacity, however, be conservative and use the 550-watt figure. Adding a monitor that draws 100 watts brings the total to 650 watts or more. To run two fully loaded AT systems, you need an 1,100-watt UPS. Don't forget two monitors, each drawing 100 watts; the total, therefore, is 1,300 watts. Using the $1 to $2 per watt figure, a UPS of at least that capacity or greater will cost from $1,300 to $2,600 dollars—expensive, but unfortunately what the best level of protection costs. Most companies can justify this type of expense for only a critical-use PC, such as a network file server.

In addition to the total available output power (wattage), several other factors can differentiate one UPS from another. The addition of a ferroresonant transformer improves a unit's power conditioning and buffering capabilities. Good units have also an inverter that produces a true sine wave output; the cheaper ones may generate a square wave. A square wave is an approximation of a sine wave with abrupt up-and-down voltage transitions. The abrupt transitions of a square wave signal are not compatible with some computer equipment power supplies. Be sure that the UPS you purchase produces a signal compatible with your computer equipment. Every unit has a specification for how long it can sustain output at the rated level. If your systems draw less than the rated level, you have some additional time. Be careful, though: most UPS systems are not designed for you to sit and compute for hours through an electrical blackout. They are designed to provide power to whatever is needed, to remain operating long enough to allow for an orderly shutdown. You pay a large amount for units that provide power for more than 15 minutes or so.

There are many sources of power protection equipment, but two of the best are Best Power and Tripp Lite. These companies sell a variety of UPS, SPS, line, and surge protectors. They are listed in Appendix B.

Using Data-Backup Systems

Making a backup of important data on a computer system is one thing that many users fail to do. A backup is similar to insurance: you need it only when you are in big trouble! Because of the cost in not only dollars, but also in time and effort, many users do not have adequate backup—which is not a problem until the day you have a catastrophe and suddenly find yourself without your important data or files. This section discusses several forms of backup hardware and software that can make the job both easier and faster, and—hopefully—cause more users to do it.

Backup is something a service technician should be aware of. After I repair a system that has suffered some kind of disk crash, I can guarantee that the disk subsystem will be completely functional. I cannot guarantee, however, that the original files are on the disks; in fact, the drive may have to be replaced. Without a backup, the system can be physically repaired, but the original data may be lost forever.

Nothing destroys someone's faith in computer technology faster than telling them that the last year or more of work (in the form of disk files) no longer exists. When I visit a customer site to do some troubleshooting or repair, I always tell my clients to back up the system before I arrive. They may be reluctant to do so at first, but it is better than paying a technician by the hour to do it. A backup must be done before I operate on a system, because I do not want to be liable if something goes wrong and data is damaged or lost. If the system is so dysfunctional that a backup cannot be performed, I make sure that the client knows that the service technician is not responsible for the data.

A good general rule is never to let a backup interval be longer than what you are willing to lose some day. You always can reload or even repurchase copies of software programs that might have been lost, but you cannot buy back your own data. Because of the extremely high value of data compared to the system itself, I have recommended for some time that service technicians become familiar with data-recovery principles and procedures. Being able to perform this valuable service gives you a fantastic edge over technicians who can only fix or replace the hardware.

Backup Policies

All users and managers of computer systems should develop a plan for regular disk backups. I recommend that one person in an office have the responsibility for performing these backups so that the job is not left undone.

A backup interval should be selected based on the amount of activity on the system. Some users find that daily backups are required, and others find that a weekly arrangement is more suitable. Backups rarely must be scheduled at more than weekly intervals. Some users settle on a mixed plan: perform weekly disk backups and daily backups of only the changed files.

The procedures for backing up and for dealing with copy protection are explained in the following sections.

Backup Procedures. You should back up to removable media such as cartridge or tape, which you remove from a system and store in a safe place. Backups performed on nonremovable media, such as another hard disk, are much more vulnerable to damage, theft, or fire; also, having multiple backups is much more expensive.

Because of the relatively low cost of hard disks, some users unfortunately install two hard disks and back up one to the other. Worse, some users split a single disk into two partitions and back up one partition to the other. These backups are false backups. If the system were subjected to a massive electrical surge or failure, the contents of both drives could be lost. If the system were stolen, again, both backups would be lost. Finally, if the system were physically damaged, such as in a fire or other mishap, both the data and the backup would be lost. These are good reasons that it is important to back up to removable media.

Perform your backups on a rotating schedule. I recommend using a tape-backup system with at least three tapes per drive, in which you back up data to the first tape the first week. The second week, you use a second tape. That way, if the second tape has been damaged, you can use the preceding week's backup tape to restore data. The third week, you should use still another, third, tape and place the first tape in a different physical location as protection against damage from fire, flood, theft, or another disaster.

The fourth week, you begin to rotate each tape so that the first (off-site) tape is used again for backup, and the second tape is moved off-site. This system always has two progressively older backups on-site, with the third backup off-site to provide for disaster insurance. Only removable media can provide this type of flexibility, and tape is one of the best forms of removable media for backup.

Dealing with Copy Protection. One thing standing in the way of proper backups of some of your software is copy protection, a system in which the original disk the software is on is modified so that it cannot be copied exactly by your system. When the programs on the disk are run, they look for this unique feature to determine whether the original disk is in the system. Some software makers force you to use master copies of their programs by using this technique to require that the original disks be placed in the floppy drive for validation even though the system might have the software loaded on the hard disk.

Some forms of copy protection load the software on a hard disk only from an original disk, and modify the hard disk loaded version so that it runs only if it remains on the system in a specified set of sectors. If the program ever is moved on the disk, it fails to operate. Because these requirements make the software highly prone to failure, copy protection has no place on software used in a business environment.

My personal response to copy protection is to refuse to use, buy, or recommend any software from a company engaged in this practice. With rare exceptions, I simply do not buy copy-protected software. Unprotected alternatives always are available for whatever type of program you want. You might even discover that the unprotected alternative is a better program. If you don't make the software-purchasing decisions in your organization, however, you may have little choice in this matter.

Experienced computer users know that you never should use original disks when you install or configure software. After I purchase a new program, I first make a copy and store the original disks. In fact, I use the original disks solely for making additional backup copies. Following this procedure protects me if I make a mistake in installing or using the software. Because of the need for backup and the fact that copy protection essentially prevents proper backup, a solution has been devised.

When you must use a piece of protected software, you can purchase special programs that enable you to back up and even to remove the protection from most copy-protected programs on the market. Two such programs are CopyWrite, by Quaid Software Ltd., and Copy II PC, by Central Point Software. These programs cost about $50 and are absolutely necessary when you are forced to deal with copy-protected software.

Note that no matter what a software license agreement says, you have a legal right to back up your software; this right is guaranteed under US copyright law. Do not let a software license agreement bamboozle you into believing otherwise.

The best way to fight copy protection is with your wallet. Most companies respond to this economic pressure; many respond by removing the protection from their programs. Fortunately, because of this economic pressure by influential users, the scourge of copy protection has been almost eliminated from the business-software marketplace. Only a few remaining programs have this unfortunate defect; I hope that the protection will be eradicated from them as well.

Backup Software. In considering how to back up your system, you should be aware of both the hardware and software options available. This section first explores the software-only options; that is, using either what you get with DOS, or some more powerful aftermarket software to back up using a floppy drive. You will learn that the aftermarket software usually offers many features and capabilities that the standard DOS BACKUP program does not have. After discussing software alternatives, this section looks at complete, dedicated hardware and software backup systems. Using specialized hardware is the best way to have an easy, effective, and safe way of backing up your system.

The DOS BACKUP Command. The most basic backup software you can use are the DOS commands BACKUP and RESTORE. Since their introduction in Version 2.0 of DOS, and up to Version 3.3, these commands have frustrated users with their bugs and other problems. The versions supplied with DOS V4.0 and higher have been greatly improved, but they still do not offer what many aftermarket products offer. Because of the way BACKUP and RESTORE use the floppy drive as the hardware device, you can use this software only for backing up systems with low-capacity hard drives.

New Backup Software with DOS 6.x. Both Microsoft and IBM have included new backup software with their respective versions of DOS 6.x. These programs far outstrip the capabilities of the original BACKUP command and are much easier and safer to use.

In MS-DOS, Microsoft supplies a limited version of the Norton Backup software called MSBACKUP. This program is a full-featured menu-driven program that is designed to back up to floppy drives only. It is very easy to use, and represents a great leap from the older BACKUP and RESTORE commands. Because it is really a restricted version of the Norton Backup program by Symantec, you can easily upgrade to the more full-featured Norton Backup program and still retain compatibility with all of your existing backups.

IBM went a different route in PC DOS and supplies an only slightly limited version of the Central Point Backup program called CPBACKUP. This offers a great deal more functionality than the MSBACKUP program supplied by Microsoft. CPBACKUP can back up to a variety of different devices, including tape drives. It is also an easy to use menu-driven program, but can also be completely automated with an extensive number of command line options. Again, a more full-featured version can be obtained directly from Central Point which would be compatible with your existing backups.

Aftermarket Software for Floppy Backups. Most aftermarket software offers convenience and performance features not found in the DOS BACKUP and RESTORE commands. If you are relegated to using the floppy disk drive as your backup hardware, you should do yourself a favor and investigate some of the aftermarket software designed for backup. I recommend FASTBACK, by Fifth Generation Systems, as well as the Norton (Symantec) and Central Point backup programs. The latter two are included in utility packages from these companies, including Norton Desktop for Windows and PC Tools.

Even with these programs, however, backing up a hard disk larger than 20 to 40 megabytes is not recommended because it requires using a large number of floppy disks. To explain, consider the system I use. My main portable system has a 1-gigabyte disk drive, which requires *728* 1.44M floppy disks for backup. Some backup programs perform data compression that can reduce that number by a third to a half; optimistically, therefore, you still would need nearly 400 high-density disks. Because I always perform backups to a rotating set of media, with three backups in the rotation, I need somewhere from *1,200* *to 2,184 HD disks for my backup*. That number handles only one of my systems; I have an aggregate of more than 1 gigabyte of disk space on several other systems as well. When you imagine trying to manage more than a thousand HD floppy disks, not to mention feeding 400 or more of them into a system for a day or so to perform each backup, you can see the problem associated with backing up large drives to floppy disks.

One solution is to back up the drive using either an 8mm videotape drive or a 4mm digital audio tape (DAT) drive. Either type of tape system easily backs up more than one gigabyte to a single tape and can perform a complete backup of a full drive in just over three hours. Partial backups take only seconds or minutes.

Another feature of a tape drive is that, because it is an external unit, I can carry it around and use it to back up all my systems. Also, the media costs for tape are much less than for floppy disks: the 4mm 1.3 gigabyte tapes cost only $20 or less. Therefore, I can have three backups that cost less than $60—using floppy media would cost from $1,000 to $2,000 or more. You may not have a gigabyte of on-line storage to back up, but any amount over 40 megabytes starts to become unwieldy when you are using a floppy drive.

For a total of about $1,500 for some of the less-expensive DAT or 8mm units, I have a reliable, complete, high-speed, and simple backup of every system I own. If you have more than 40M to back up—or more than one system—a tape-backup device is the best way. Various tape-backup products are discussed in the following section.

Dedicated Backup Hardware

As explained earlier, you can choose from software and hardware options when you consider backup equipment. In this section, you learn about dedicated backup hardware, usually a tape or cartridge device designed for high-speed, high-capacity backup purposes.

Backup Systems. A good, reliable backup is important when you're using a large hard disk. With a disk of 40M or more, you should consider some form of hardware backup device other than only the floppy drive. Tape backup is available in configurations easily supporting hundreds of megabytes and more. Tape backup is fast and accurate.

Because your data is probably worth much more to you than the physical hardware on which it is stored, it is important to develop a backup system you will use. A tape system makes the backup convenient and, therefore, helps to ensure that you will complete it. If backing up your system is a difficult, time-consuming operation, you or the person responsible for doing it probably will not do it regularly (if ever). Tape units also can be set to perform unattended automatic backups during times when you are not using your computer, such as during the middle of the night.

Basic parameters of different tape-backup devices include the following:

- Type of media used
- Hardware interface
- Backup software

These parameters are explained in the following sections.

QIC Standards. The Quarter-Inch Committee (QIC) issues standards for such things as tape-drive controller adapters, the tape cartridges, and commands that tape drives understand. You will find them listed in Appendix B, the "Vendor List."

The QIC organization was formed to develop standard formats for DC-600 and DC-2000 data cartridge tape drives. Because quarter-inch tapes typically have been the market leaders, even suppliers of other systems generally follow QIC standards, such as the SCSI command set in QIC-104, and QIC-122 or QIC-123, which define data-compression methods. Data compression allows you to store more data on a certain length of tape. Vendors can change their backup software so that the codes are inserted at the beginning of the tape. Then, if the tape is read back by other tape drives, the read back drive can determine the method used.

Tape Media. I like to use systems that employ industry-standard media. *Media* refers to the type of tape format used. The type of media you select dictates the capacity of the tape-unit storage. Many different types of systems use many different types of media, but I am concerned with only a handful of standard media types. Four primary standards are the 3M Data Cartridge 600 (or DC-600), the DC-2000, the 4mm digital audio tape (DAT), and the 8mm (video) cartridge media. I recommend any of these types, depending on your requirements for capacity and performance. This list shows the different media types:

- *DC-600 cartridges*. Invented by 3M, the DC-600 cartridge is a relatively large tape, 4 × 6 × 0.665 in L × W × D, and has a heavy metal base plate. This long-running standard was introduced in 1971. Many variations of the DC-600 tapes exist; longer lengths have capacities anywhere from 60M, at the low end, to 525M or more at the high end.
- *DC-2000 cartridges*. Also invented by 3M, this cartridge is one of the most popular media for backup purposes. This tape is 2.415 × 3.188 × 0.570 inches in L × W × D, with a heavy metal base plate. Although early versions held only 20M, the most common ones now hold 80M or more. A variety of capacities result from the different tape formats available.
- *4mm DAT (digital audio tape) cartridges*. At first glance, this cartridge looks like a regular audiocassette but is slightly smaller. These cartridges come in several recording formats. The most widely used is digital data storage, or DDS for short. Another format, DataDAT, is also available. The recording technology is similar to digital audio tape decks and is done in a Helical-Scan format; it is licensed by Sony (the original DAT developer). DAT tapes can hold as much as 1.3 gigabytes (1 gigabyte = 1,024M).
- *8mm cartridges*. These cartridges, which use the 8mm videotape developed by Sony for camcorders, comprise one of the highest-capacity backup systems available. Most units store 2.3G but larger-capacity units are available. The recording is done physically as a Helical-Scan, the same method used in a video recorder.

The DC-600 drives now store from 60M to 525M or more, depending on the format and quality of tape used. The DC-2000 media stores only 40M to 80M or more per tape and might be suitable for low-end (such as home) use. Because the larger-capacity 4mm and 8mm units store more than 1 or 2 gigabytes on a tape, they are ideal for network server applications. These large-capacity units can cost from $1,500 to $6,000 or more, depending on the features included and—of course—where you buy it.

Make sure that the system is capable of handling your largest drive either directly or through the use of multiple tapes. With this unit, you can change tapes in the middle of a backup session to accommodate greater capacities; because this method is inconvenient and requires an operator, however, you eliminate unattended, overnight backups. For network applications or large storage requirements, or anything else for which you need top performance and the most reliable backup, you should select DC-600 or the 4mm or 8mm rather than the DC-2000 media.

I recommend only tape systems that use these media because they offer the greatest value per dollar, are the most reliable, and hold a large amount of data. Stay away from devices that use the 3M DC-1000 tapes, Phillips audio tapes, or VHS VCR tape systems. These systems often are slow, error prone, and inconvenient to use; some do not handle larger capacities, and they are not standard.

Interfaces for Tape Backup Systems. Of the three primary hardware-interface standards, virtually all professional backup systems use the QIC-02 (for Quarter-Inch Committee 02) or SCSI (for *Small Computer Systems Interface*), and more recently, the parallel port interface. Some low-end systems use the Shugart Associates 400 (SA-400) interface, which is the standard floppy controller. The interface you select controls the backup speed, whether an adapter card and slot are required, and the reliability and capacity of the unit.

In the PC and XT systems with their 4-drive floppy controllers, you can use the extra connector on the back of the existing floppy controller for some tape-backup units. This connector saves the use of a slot in these systems. Unfortunately, the AT systems require some sort of multiplexer card to enable sharing one of the internal floppy ports or another complete floppy controller so that these systems can work.

I don't recommend these floppy interface systems in general because they are slow, and the floppy controller lacks any form of error-detection and correction capability, which can make the backups unreliable. A system using the SA-400 interface has a maximum data rate of 2M per minute, less than half the QIC-02 or SCSI interfaces. These floppy interface systems are suitable for home computer users on a tight budget but should not be used in business or professional environments. Another problem is that the interface limits the capacity to 40M, and you might have to format the tapes before using them.

The QIC-02 interface, designed specifically for tape-backup products, represents an industry standard. Many companies offer products that use this interface. The QIC-02 can back up at a rate of 5M per minute. The interface usually is a short adapter card that requires a free slot in the system. You can buy these adapters separately and enable many systems to use the same externally mounted drive. A free slot is required, however, and might not be available with a regular PC.

Perhaps the best hardware interface is the Small Computer Systems Interface (SCSI), which enables your PC to back up at high data rates. The data throughput of the SCSI interface is either 5MB or 10MB per minute for standard and fast 8-bit SCSI respectively. There are currently no tape drives or hard disks that can sustain this rate of transfer, so the performance is dictated by the drives rather than the SCSI interface. Because of the high speed of the SCSI interface relative to the tape and hard disk drives, when faster media and drives are available, SCSI can easily take advantage of the greater speed.

The SCSI interface is supplied as a SCSI host adapter card that plugs into the system unit and can connect to the tape drive. The SCSI host adapter requires an expansion slot, but it can connect to other devices such as hard disk drives or tape drives with embedded SCSI interfaces. The use of a single host adapter for as many as six hard disk drives and one tape unit is a strong case in favor of SCSI as a general disk and tape interface. The only real problem with SCSI is that you will need drivers which support both your particular tape drive as well as your particular host adapter. Most tape drivers support host adapters that have ASPI (Advanced SCSI Programming Interface) or CAM (Common Access Method) drivers. In other words, you will need the ASPI or CAM driver for your host adapter, and a tape driver that connects to ASPI or CAM for your tape drive. Although there can sometimes be problems getting the correct drivers for a particular system, SCSI is still the highest performance and most flexible interface for tape backup systems.

Another interface gaining wide acceptance is the parallel port. It is used for more than just printers, as several companies have introduced versions of their backup hardware and software systems that use the parallel port to connect to your system. The parallel port offers a high speed data transfer interface with easy configuration and cabling. Unlike serial ports, there is usually no confusion as to how the cables are wired for parallel port connections.

Several companies, like MicroSolutions and Colorado, offer complete external backup systems including all hardware, software, and cables needed to connect to the parallel port of any system. These types of backup systems offer independence from a particular type of system, which is especially useful in a mixed machine or mixed bus environment.

Other companies, like Trantor, offer parallel port SCSI adapters, to which you can connect any SCSI device including hard disk, CD-ROM, tape backup, scanners, and so on. If you are out of slots or desire to set up an external backup system that will attach to and back up several systems, I recommend using one of these parallel port SCSI adapters like the Trantor MiniSCSI Plus. Using a parallel port SCSI adapter and one of the higher performance DAT tape drives, you can set up a highly capable external backup system that can function on virtually any machine.

Tape Backup Drivers and Software. Now you need to consider the software that will run the tape system. Some manufacturers have written their own software, which is proprietary and used only by that manufacturer. In other words, even though you might be using a tape unit with the same media and interface, if the software is different you cannot interchange data between the units. One type of software does not recognize the data formats of another proprietary software system. Other tape software packages are designed to run on a variety of tape hardware rather than a single manufacturer's system. Several hardware independent backup software packages support a variety of tape drives such as SyTOS (Sytron Tape Operating System) by Sytron, Novaback by Novastor, and Central Point Backup by Central Point. All of these products work with a number of different tape drives. SyTOS is somewhat of an industry standard because it is endorsed (and sold) by IBM and many other vendors, as well as having versions that work on all major PC operating systems.

SyTOS has been selected by IBM as the standard software for its units. Because of market pressures to be compatible with IBM, Compaq and other system and tape-drive vendors are offering SyTOS with their tape systems. Because two systems that use the same media, interface, and software can read and write each other's tapes, you can have data interchangeability among different tape-unit manufacturers. You should consider using a system that runs SyTOS software for those reasons. Of course, the proprietary software is generally just as good, if not better. You can exchange data tapes only with other users of the same systems.

All software should be capable of certain basic operations. Make sure that your software does what you want, and consider buying only software that has these features:

- Can back up an entire DOS partition (full volume backup)
- Can back up any or all files individually (file-by-file backup)
- Allows a selective file-by-file restore from a volume backup as well as a file-by-file backup
- Can combine several backups on a single tape
- Can run the software as commands from DOS BATCH files
- Works under a network
- Can span a large drive on multiple tapes
- Can be completely verified

If the software you are considering does not have any of these essential features, look for another system.

Physical Location: Internal or External. Physical location of the tape-backup units is a simple factor often not considered in detail. I almost never recommend any tape-backup unit mounted internally in a system unit. I recommend that you buy a tape unit externally mounted in its own chassis and one that connects to the system unit through a cable and connectors.

Tape units are relatively expensive, and I never want to tie up a backup unit for only one system. The amount of time that just one system spends using the unit makes this proposition wasteful. With an external unit, many systems can share the backup unit. I just equip every system unit with the required interface, if they aren't already equipped. Extra QIC-02 cards are available for about $100 each, and SCSI host adapters are about $200. Note that many systems now include a SCSI host adapter as part of the motherboard, or as a card already installed. I already share my system among more than five computers and will continue adding systems to the backup pool. In some companies, the backup unit is mounted on a wheeled cart so that workers easily can move it from computer to computer.

If you have only one system or if all your data is stored in a single file-server system, you might have a legitimate argument for using an internal tape drive. You might wish that you had an external system, however, on the day a new system arrives, or when the server is down and you have to get a new one up and running.

Recommended Backup System

Although a variety of backup systems are available, I recommend units from any manufacturer that meets these requirements:

Type of media: DC-2000, DC-600, 4mm DAT, or 8mm

Hardware interface: QIC-02, Parallel Port or SCSI

My preferences lean towards the 4mm DAT drives because of their high capacity (4G plus with built-in data compression) and relatively high performance. Also, the media cost is very low, around $15 per cartridge. DAT devices almost always use a SCSI interface, which means you would need a SCSI host adapter installed in the system. This could be either a standard type of SCSI host adapter in a slot (or integrated into the motherboard of some systems), or for a lower cost portable solution, you could run the SCSI DAT drive from a parallel port with one of the parallel to SCSI adapters like those from Trantor. The parallel port adapters make it easy to use a single external DAT drive to back up a variety of systems.

Purchasing Warranty and Service Contracts

Extended warranties are a recent trend in the computer industry. With the current fierce competition among hardware vendors, a good warranty is one way for a specific manufacturer to stand out from the crowd. Although most companies offer a one year warranty on their systems, others offer longer warranty periods, such as two years or more.

In addition to extended-length warranties, some manufacturers offer free or nearly free on-site service during the warranty period. Many highly competitive mail-order outfits offer service such as this for little or no extra cost. Even IBM has succumbed to market pressure to lower service costs, and has an option for converting the standard one year warranty into a full-blown, on-site service contract for only $40. This same option can be extended to cover monitors and printers for only a small additional cost.

Tip

IBM and other companies are beginning to offer extended-length warranties and free or low-cost on-site service. IBM has a somewhat unknown option to upgrade the standard one year warranties on its PS/2 systems into a full-blown, on-site service contract. Under this option, you can upgrade the standard Customer Carry-in Repair warranty to an IBM On-site Repair warranty for only $40 for the first year.

If you have only the carry-in warranty, as its name implies you must carry in your computer to a depot for service; with the on-site warranty, the service technician comes to you. One car trip to a service center costs me more than $40 in wasted gasoline and time. I gladly pay the fee to have the service technician come to see me. Be sure to ask for this option when you buy your system. If your dealer is unfamiliar with it, ask that your order include IBM feature code #9805.

In most normal cases, service contracts are not worth the price. In the retail computer environment, a service contract is often a way for a dealer or vendor to add income to a sale. Most annual service contracts add 10 to 15 percent of the cost of the system. A service contract for a $5,000 system, for example, *costs $500 to $750 per year*. Salespeople in most organizations are trained to vigorously sell service contracts. Much like in the automobile sales business, these contracts are largely unnecessary except in special situations.

The high prices of service contracts also might affect the quality of service you receive. Technicians could try to make their work seem more complex than it actually is to make you believe that the contract's price is justified. For example, a service technician might replace your hard disk or entire motherboard with a spare when all you need is low-level formatting for the hard disk or a simple fix for the motherboard such as a single memory chip. A "defective" drive, for example, probably is just returned to the shop for low-level formatting. Eventually, it ends up in somebody else's system. Replacing a part is faster and leaves the impression that your expensive service contract is worth the price because you get a "new" part. You might be much less impressed with your expensive service contract if the service people visit, do a simple troubleshooting procedure, and then replace a single $2 memory chip or spend 15 minutes reformatting the hard disk.

With some basic troubleshooting skills, some simple tools, and a few spare parts, you can eliminate the need for most of these expensive service contracts. Unfortunately, some companies practice deceptive servicing procedures to justify the expensive service contracts they offer. Users are made to believe that these types of component failures are the norm, and they have a mistaken impression about the overall reliability of today's systems.

Tip

If you have many systems, you can justify carrying a spare-parts inventory, which can also eliminate the need for a service contract. For less than what a service contract costs for five to ten systems, you often can buy a complete spare system each year. Protecting yourself with extra equipment rather than service contracts is practical if you have more than ten computers of the same make or model. For extremely time-sensitive applications, you might be wise to buy a second system along with the primary unit—such as in a network file-server application. Only you can make the appropriate cost-justification analysis to decide whether you need a service contract or a spare system.

In some instances, buying a service contract can be justified and beneficial. If you have a system that must function at all times and is so expensive that you cannot buy a complete spare system, or for a system in a remote location far away from a centralized service operation, you might be wise to invest in a good service contract that provides timely repairs. Before contracting for service, you should consider your options carefully. These sources either supply or authorize service contracts:

- Manufacturers
- Dealers or vendors
- Third parties

Although most users take the manufacturer or dealer service, sometimes a third-party tries harder to close the deal; for example, it sometimes includes all the equipment installed, even aftermarket items the dealers or manufacturers don't offer. In other cases, a manufacturer might not have its own service organization; instead, it makes a deal with a major third-party nationwide service company to provide authorized service.

After you select an organization, several levels of service often are available. Starting with the most expensive, these levels of service typically include:

- Four-hour on-site response
- Next-day on-site response
- Courier service (a service company picks up and returns a unit)
- Carry-in, or "depot," service

The actual menu varies from manufacturer to manufacturer. For example, IBM offers only a full 24-hours-a-day, 7-days-a-week, on-site service contract. IBM claims that a technician is dispatched usually within four hours of your call. For older systems, but not the PS/2, IBM also offers a courier or carry-in service contract. Warranty work, normally a customer carry-in depot arrangement, can be upgraded to a full on-site contract for only $40. After the first-year $40 contract upgrade expires, you can continue the full on-site service contract for standard rates. Table 18.2 lists the rates for IBM service contracts after the warranty has expired.

Table 18.2 PS/2 Service-Contract Fees after Warranty for Annual IBM On-Site Repair (IOR) Service

PS/2 Model	Contract	PS/2 Model	Contract
8525-x01/x02	$95	8570-0x1/1x1	$450
8525-x04/x05	$110	8570-Axx	595
8525-x06/x36	$210	8570-Bx1	645
8530-001	$190	8573-031/061	385
8530-002	$140	8573-121	385
8530-021	$190	8573-161	750
8530-E01/E21	$190	8573-401	850
8530-E31/E41	$190	8580-041	360
8535-all	$260	8580-071	425
8540-all	$400	8580-081	550
8543-044	$430	8580-111/121	500
8550-021	$220	8580-161	600
8550-031/061	$220	8580-311	575
8555-0x1	$260	8580-321	605
8555-LT0/LE0	$260	8580-A21	675
8557-045/049	$400	8580-A16/A31	785
8560-041	$310	8590-0G5/0G9	800
8560-071	$345	8590-0J5/0J9	850
8565-061	$425	8590-0KD	950
8565-121	$475	8595-0G9/0GF	1,400
8565-321	$550	8595-0Jx	1,450
8570-E61	$415	8595-0Kx	1,550

If you have bought a service contract in the past, these prices might surprise you. With the PS/2 systems, IBM has rewritten the rules for PC servicing. The same type of on-site annual contract for an earlier 20M AT system, for example, costs nearly $600 annually, compared with $220 for the PS/2 Model 50. Smaller third-party service companies are having difficulty competing with these newer prices. IBM claims that the PS/2 systems are five times more reliable than the earlier systems, and the service-contract pricing is about one-third what the earlier systems cost. If the claims of additional reliability are true, IBM is doing well even with lower pricing. If the claims are not true (which is unlikely), then IBM would be losing money on its service contracts.

Tip

In summary, for most standard systems, a service contract beyond what is included with the original warranty is probably a waste of money. For other systems that have not yet achieved commodity status, or for systems that must be up and running at all times, you might want to investigate a service-contract option, even though you might be fully qualified and capable of servicing the system.

The PS/2 systems are a somewhat special case; they have low-priced, on-site warranty upgrades and fairly inexpensive service contracts compared to the costs of the systems. Remember, however, that you can buy most, if not all, parts for these systems from non-IBM sources; also, third-party companies that specialize in difficult items such as motherboards and power supplies can repair these items. Only after carefully weighing every option and cost involved can you decide how your systems should be serviced.

Summary

This chapter has presented the steps you can take to ensure proper operation of your system. It has examined active and passive preventive maintenance—the key to a system that gives many years of trouble-free service. You have learned about the procedures involved in preventive maintenance and the frequency with which these procedures should be performed.

Backup was discussed as a way to be prepared when things go wrong. The only guarantee for being able to retrieve data is to back it up. In this chapter, you also have learned about backup options.

Finally, you have learned about the commonly available warranty and service contracts provided by computer manufacturers. Sometimes the contracts can save you from worrying about tough-to-service systems or systems whose parts are largely unavailable on short notice.