



The Hidden Security Risk: Hard Drive Remanence Data


Protecting Sensitive Information with Effective and Secure Hard Drive Cleansing

Pinion Software, Inc.

July, 2004

This document is proprietary, company confidential, and access restricted. It may not be copied except with written permission of Pinion Software, Inc. The document is Copyright© 2005, Pinion Software. Pinion™ and the Pinion logo are trademarks of Pinion Software. All rights are reserved. All other trademarks, registered trademarks, service marks, and registered service marks are property of their respective owners.

The computer industry has thrived on technological advancement since its birth. It has followed the general trend defined by Moore's law, that the performance of PC's will double every 18 months. This means that after only three years a personal computer is only 25% as efficient as currently available models. As a result PC's are being made obsolete within three years. But what happens after they have been taken out of service?



It is estimated that over 300 million computers will be retired this year. In addition, more than 70% of U.S. companies lease computers that must be returned at the end of the lease term.

Companies purchasing computers have several options for passing on old computers that include taking it to a PC recycler or toxic waste disposal center, trading it in towards a new system, selling it, or donating it to a worthy charity. As the personal computer has risen in popularity, the need for dealing with all the old hardware accumulating in homes and businesses has become a significant problem. Filling up landfills with large computers that are only a few of years old is not a good idea nor environmentally friendly, especially considering the foul substances that can leech out of old computers.

In the year 2000, the National Recycling Coalition, an industry group, conducted the first large-scale survey of America's electronic recycling efforts. The results were staggering. Between now and 2007, the survey found 500 million personal computers will become obsolete.

About 150,000 hard drives were "retired" last year, the research firm Gartner Dataquest estimates. Many ended up in trash heaps, but many also found their way to secondary markets.

With increasing pressure to reduce costs and the availability of new methods to resell computers, businesses are looking for ways to either internally recycle their aging computer inventory or sell them into a growing used computer market. It is not unusual to find companies reselling their excess equipment on Internet sites such as eBay. However, in all cases there is a need to remove all data stored on the computer before its disposal.

The loss of confidential information left on resold PCs can be catastrophic in today's information rich economy.

1 Data Storage Basics

The first part to understanding the scope of the data removal problem is to understand the basics of data storage within the computer system. There are fundamentally two ways of retaining data in the PC, RAM memory and disk. In the case of the disk storage the principle device is the hard drive. What is not commonly known is how these systems work and how they interact.

The hard drive is used to store both program information and data so that they can easily be retrieved. The data can be generated by the end user or the programs themselves. The not so obvious part of the previous statement is the way in which the programs use the hard drive. One of the major challenges with many operating systems is that they cannot hold all of the data that they need in the system RAM memory. So the operating system will periodically "swap out" the data from RAM memory onto the disk. This is known as a swap file. A lot of information is contained in the swap file, such as the last email you read or the last spread sheet you used. The applications themselves also use the hard drive to store data. Ever wondered just how an application can recover a damaged file or reverse a change. Copies of the data are stored on the hard drive as back-up.

Before a hard drive can be used it has to go through some steps to prepare it to accept information so that it can be easily retrieved. This occurs in two steps. The first is to establish the areas on the drive and how they are going to be used. A common utility for this is FDISK. This will establish the sections, or partitions, on the hard drive and let the operating system know which one will contain the operating system program. The next step is to format the drive. This sets up an environment on the disk so that the operating system can store and access files from the drive.

2 Myths about Data Removal

Myth #1 – I can just empty my recycle bin

One of the good and bad things about the standard Microsoft Windows operating system is that it has been built to be efficient. This leads to some serious misunderstandings on how data is removed from a hard drive. As many users will already know, when a file is deleted by a Windows delete command, it is not really removed; it just goes to the Recycle Bin. But once the recycle bin is emptied it is gone. Unfortunately, no it isn't. The operating system does not really remove the data. What it does do is make it difficult to find. There are many ways to perform an "undelete" on the data that has been processed when the Recycle Bin has been emptied.

Myth #2 - I can reformat the drive

This myth falls into the same category as the previous myth. When the drive is reformatted the utility will merely rewrite the information that is used to locate the files on the drive.

Over two years, Simson Garfinkel and Abhi Shelat assembled a collection of 158 used hard drives, shelling out between \$5 and \$30 for each drive at secondhand computer stores and on eBay.

Of the 129 drives that functioned, 69 still had recoverable files on them and 49 contained "significant personal information" -- medical correspondence, love letters, pornography and 5,000 credit card numbers. One even had a year's worth of transactions with account numbers from an automated teller machine in Illinois.

Even formatting a drive may not do it. Fifty-one of the 129 working drives the authors acquired had been formatted, but 19 of them still contained recoverable data.

Essentially it will tell the operating system that there are no files and that all of the space on the disk is free. Until the operating system comes along and writes new data over the old, the original data is still there. There are tools available, many are free, that will perform an un-format on a previously formatted drive and retrieve any previous data that has been left on the drive.

Myth #3 - I can just run FDISK on the drive again

This is probably one of the most pervasive myths. Again, as in the previous two myths, the data is not removed. In the case of an FDISK operation, all of the information that is needed to locate the data from the operating system is removed. But as in the reformatting case, the original data is still there in its raw forms. The data can be recovered by examining the binary information on the drive. Now this may not sound very useful until it is realized that there are now many programs that will automatically convert the binary data into text. Furthermore, full text searches can be performed on the drive after it has been through an FDISK cycle. With these tools large portions of data can be extracted even though the disk is presumed "clean".

The Bottom Line

None of the standard tools described above will remove the bulk of the data contained on the hard drive. The only solution to ensure that the information on the hard drive is removed is to either physically destroy the drive itself, or write over the existing data so that it cannot be recovered

3 US Department of Defense 5220.22-M standard.

There has been a standard in place for some time that addresses the problem of permanent removal of data from a hard drive. The standard was developed by the Defense Security Service (DSS) and is used by many federal and commercial organizations. As a Department of Defense (DoD) agency, DSS makes its contribution to the national security community by conducting personnel security investigations and providing industrial security products and services, as well as offering comprehensive security education and training to DoD and other government entities. DSS administers three industrial security programs, of which the National Industrial Security Program (NISP) is the largest.

Under the NISP, DSS Industrial Security Representatives oversee cleared contractor facilities and assist the organizations' management staff and Facility Security Officers in formulating their security programs. As part of the NISP initiative, DSS has developed the DoD standard 5220.22-M NATIONAL INDUSTRIAL SECURITY PROGRAM OPERATING MANUAL. Among other items, the standard outlines the method that is to be used for removing data from unclassified hard drives – sanitizing. The standard defines Sanitization as follows;

Sanitization. Sanitization is the process of removing the data from media before reusing the media in an environment that does not provide an acceptable level of protection for the data that was in the media before sanitizing. IS resources shall be sanitized before they are released from classified information controls or released for use at a lower classification level.

NISP defines an overwriting technique that will remove any existing data yet leave the hard drive in a state where it can be reused. The process involves the following two steps:

1. Before any sanitization product is acquired, careful analysis to the overall costs associated with overwrite/sanitization should be made. Depending on the contractor's environment, the size of the drive and the differences in the individual products time to perform the sanitization, destruction of the media might be the preferred (i.e., economical) sanitization method.
2. Overwrite all addressable locations with a character, then its complement. Verify "complement" character was written successfully to all addressable locations, then overwrite all addressable locations with random characters; or verify third overwrite of random characters. Overwrite utility must write/read to "growth" defect list/sectors or disk must be mapped before initial classified use and remapped before

sanitization. Difference in the comparison lists must be discussed with the DSS Industrial Security Representative (IS Rep) and/or Information System Security Professional (ISSP) before declassification. Note: Overwrite utilities must be authorized by DSS before use.

The full matrix of recommended disposal methodologies for a wide variety of computer components is available at http://www.dss.mil/infoas/clearing_and_sanitization_matrix.doc which is part of the Information Assurance Web Site. The site is hosted by the Defense Security Service (DSS) Industrial Security Information Assurance Branch which comprises of computer security specialists and computer scientists who support existing Industrial Security programs. The site has been designed to help you with all Information System (IS) issues as they relate to the National Industrial Security Program (NISP) and can be found at <http://www.dss.mil/infoas/index.htm>.

4 Other considerations when choosing a sanitizing product

In addition to meeting the process defined by the DoD 5220.22-M standard there are some other important requirements that should be taken into consideration before selecting a product.

BIOS Independence

When preparing a hard drive for disposal or recycling, it is important to remove all of the program files as well as the data files. If the program files are left, not only is the software shipped with the disk, but all of the data files that have been generated by the operating system, such as swap file and temporary files, are typically still available. To effectively sanitize the operating system, the sanitizing application must function independently. The sanitizing application may still use a permanent program incorporated in the PC hardware that allows the OS to communicate with the various hardware devices. BIOS (basic input/output system) is the program a personal computer uses to get the computer system started after you turn it on. It also manages data flow between the computer's operating system and attached devices such as the hard disk, video adapter, keyboard, mouse, and printer.

In many cases the BIOS used can be several years old. This can cause problems when a sanitizing application uses the BIOS to acquire the hard drive size. Older BIOS's can return an incorrect value when it is not compatible with a newer larger hard drive. This is not noticed during normal operation as the flaw is automatically corrected by the operating system. However if the sanitizing product does not correct the flaw, then it will only remove the data from part of the hard drive. This will result in data being left behind on the disk. Only a BIOS independent sanitizing product will obtain the correct drive size and completely remove all of the data.

Hard Drive standard compatibility

There are two predominant standards for hard drive technology used by personal computers today. One is IDE (Integrated Drive Electronics) which is based on the IBM PC Industry Standard Architecture (ISA). Most computers sold today use an improved version of IDE called Enhanced Integrated Drive Electronics (EIDE). The other is SCSI (pronounced SKUH-zee and sometimes colloquially known as "scuzzy"), the Small Computer System Interface, is a set of ANSI standard electronic interfaces that allow personal computers to communicate with hard drives. The sanitizing utility should be able to sanitize either drive type.

Size compatibility

As hard drive sizes continue to increase it is important to verify that the sanitizing product is able to address the larger size of drives. Hard drive sizes have already exceeded the 100 gigabyte limit. Many products are not yet capable of handling this size of drive.

Reporting

An important part of any business process is accounting. This is also true for sanitizing a hard drive for two reasons. First there needs to be a record that all of the software that was on the drive has been removed. This will allow the software to be legally re-used on another computer. Second; by having a record that all company information has been removed the drive can then be resold outside of the company. This can be especially true if there are regulatory restrictions on how businesses must protect the confidentiality of any information that they have received.



5 Summary

As computer systems become faster and cheaper, the desire to replace them in the workplace will result in the need to dispose of the obsolete equipment. Although this equipment may not meet the needs of the business there is a thriving market, especially for personal use, for reselling it. However it is important that no digital property is lost in this transaction. If this occurs the impact can range from inconvenience, public embarrassment, fiscal damage or violations of regulatory standards. The DoD standard 2550.22-M provides a good, proven framework for designing a digital data disposal process. This can be augmented by some other considerations that are not currently included in the standard to help select the right sanitizing product. This will result in meeting the goal of retiring obsolete equipment and recovering any residual value while not compromising digital data security.

6 Contact Information

Mailing Address	Pinion Software, Inc. 4030 West Braker Lane, Suite 450 Austin, Texas 78759
Sales	sales@pinionsoftware.com (800) 308-5825
Technical Support	support@pinionsoftware.com voice: (512) 583-0868