

Counter forensics techniques – a brief overview

Noemi Kuncik and Andrew Harbison

This article was originally published in Digital Forensics magazine

Introduction

It is well known that computer evidence can easily be changed and just as easily deleted. Over the years, in the analysis of computer based evidence we would very occasionally come across cases where an individual altered documents in an attempt to confuse or mislead our clients, or had deleted large numbers of files in an attempt to prevent them reading them.

These attempts were rarely much of a problem. IT forensics tools are very effective at identifying altered documents and retrieving deleted data when the alteration and deletion is done using conventional tools.

Up to a couple of years ago we would rarely see sophisticated attempts to dispose of evidence, such as “file shredding” and “evidence eliminating” tools. This activity was found only when investigating highly skilled computer specialists.

In the last couple of years things have begun to change. Electronic discovery and IT forensic support of litigation have become far more common together with the use of sophisticated “counter forensic” techniques to inhibit the recovery of valid evidence from computers. It is now routine procedure to check for data tampering in every case. Correspondingly IT forensics specialists now have to acquire a new set of skills, while at the same time becoming less heavily reliant on the standard forensic applications, some of which counter forensic tools are designed specifically to subvert.

What kinds of counter forensic techniques are there?

There are four broad options open to an individual trying to prevent or inhibit the investigation and analysis of data on a computer.¹ They can simply attempt to destroy the data, alter it, hide it inside a computer system or they can try to pre-empt it, preventing it from accumulating in the first place.

1 Data destruction

At first glance it would appear that the best solution to dealing with incriminating data is simply to destroy it. Many individuals employ different techniques to fully remove accumulated data

¹ Harris R., *Arriving at an anti-forensics consensus: Examining how to define and control the anti-forensics problem*
Digital Investigation 3S (2006) S44 – S49

from their computers, from simply deleting it conventionally, to using sophisticated “evidence elimination” software, to actually replacing the evidential hard drive.

Among the data removal techniques available are:

- **File deletion.** The biggest problem with file deletion is that computers store information in many different locations and, in most cases, simple file deletion will do practically nothing to remove it.

The computer will have created link files and other “tags” in the operating systems as well as references in the registry. All these will indicate that files that once were stored on the computer are no longer there. Also, because many computer applications make temporary copies of the files on which they work, deletion is no guarantee that the data stored in a file cannot be recovered.

- **Reformatting.** Many people think that reformatting a drive will destroy everything on it, largely because when they do it the process sometimes takes some hours to complete. In reality, it could easily take a forensics investigator less time to reverse the effects of a disk reformatting than it does to carry it out in the first place.

The main action carried out by a format is resetting the hard drive’s file table. This is equivalent to removing all the index cards from a library indexing system. It does nothing to the files on the disk. The format will also look for defects on the surface of the disk (so-called “bad sectors”) which it records. This sweep for bad sectors is time consuming and can cause a format to take hours on a large hard drive. Without this sweep, reformatting will normally take a matter of seconds.²

- To reverse the format all the investigator needs to do is locate the deleted file table and reconstruct it. This is something a lot of forensics tools allow investigators to do straightforwardly.

Occasionally individuals reinstall the operating system after a format. This can cause problems because it usually overwrites the old file table. Nevertheless, deleted files on the hard drive can still be found, although the process is typically a lot more difficult and time-consuming.

- **Defragmentation.** When a computer hard drive gets very full, it becomes difficult for the computer to store large files on it. Sometimes it is impossible to store a file in one contiguous space on the drive. Instead the computer will store parts of the file in a number of different locations. The file is said to have been “fragmented”.

Fragmentation tends to slow down the computer a great deal. It means that the hard drive has to be searched in a number of different locations to assemble a file before it can be loaded in memory. This is particularly problematic because computers routinely use a number of system files which can find themselves fragmented in this way.

² For example, a full duration format of a 60GB 2.5 inch 5,400 rpm SATA hard drive, took around 27 minutes on a recently purchased NT-based forensic workstation. A “quick” format of an identical model disk took around 20 seconds on the same machine. (Good info – this helps make the point.)

Defragmentation reorganises the hard drive so that all parts of all files are stored in a single location in contiguous fashion and also concatenates all files on a computer into a single logical area on the disk allowing for a faster file search on the disk. In doing this the computer rewrites and erases files all over the disk causing disruption to data in the unallocated spaces of the hard drive.

These actions are all normal parts of the defragmentation process and are not normally problematic, unless the computer in question is under forensic analysis. Anything that disrupts the unallocated space is very likely to destroy evidential remnants written in those locations. When combined with file deletion it can greatly increase the chance that forensic traces of a deleted file are rendered irrecoverable.

As with file deletion, defragmentation is most likely to be effective in destroying evidence when the disk is nearly full or when it is performed a considerable time before analysis occurs. As with file deletion, defragmentation will still leave a lot of trace evidence across an evidential hard drive which will prove that files have been removed. Hence it is of limited effectiveness as a counter forensic technique.

Using a defragmenter is atypical behaviour and so most courts will become suspicious on a defendant who suddenly becomes enthusiastic about defragmenting their computer only when it becomes the likely subject of forensic analysis.

- **File “shredding”.** By performing file shredding, a file is not merely deleted but overwritten. The data bytes of the file stored on the disk are overwritten with new data, and in most cases the file table entry is also overwritten. File shredders are readily available both for purchase and as free downloads on the Internet.

File shredding is more effective than conventional deletion as some specific data may not be retrieved by forensic analysts. However, it retains many of the other drawbacks of conventional file deletion.³ It usually does not remove all traces of the erased files. A good investigator should be able to establish which files were on a computer and when they were removed. Hence it is a technique that still holds considerable risks for any defendant who decides to use it.

File shredding can however be used for legitimate reasons such as additional security on a computer for anyone needing to routinely destroy sensitive data.

- **Evidence elimination.** “Evidence eliminators”, named after an early example of the type, are software tools that explicitly attempt to remove as much residual data from a computer as possible that might be of interest to a forensic investigator. They are far more effective than any of the techniques already discussed in removing data.

³ For example, if a Word Document is shredded in Windows, it might leave link files in the Recent folder, MRU records in multiple locations in the Registry and fragments of any temporary files created during editing (which will themselves contain extensive metadata), and possibly other material in the Pagefile and Hiberfile.

- Almost all eliminators will include a file shredder, but many will also contain functions that will overwrite empty slack and unallocated spaces, history files, log files and registry settings. The better ones are run from a CD or USB device so as not to leave traces of themselves on the media being “cleansed”. Most of the more effective evidence eliminating applications are commercial, and some of the best can be quite expensive.

Although these methods are more successful, they are far from foolproof and can often rebound on their users. Their principal problem is that they tend to remove too much data. A competent IT forensic analyst will know to expect to find residual data in some locations on any hard drive such as some file data in the “slack” space and some file fragments in the unallocated space. If this material is not present an investigator will probably “smell a rat”. With little work a good analyst may be able to determine not only that an evidence eliminator has been employed, but when and by whom. Sadly for evidence tamperers there are few, if any, perfect evidence eliminating applications, so incriminating material can still be left behind.⁴

- **Disk wiping.** In some circumstances a defendant decides that it is worth their while to destroy evidence regardless of how much upset it might cause. There are plenty of disk wiping tools on the market, many of them free to download and most are straightforward to use. The use of disk wiping tools will represent a serious contempt of court, if undertaken after the subpoena or other notification has been received, and the penalties might be severe.

We occasionally find suspects wiping their disks and reinstalling the operating system, or “re-imaging” their computers – copying the entire contents of another disk drive onto the evidence drive – to disguise the fact that the drive has been erased. The results of such activity are so obvious, however, that it would take a very careless analyst to miss them. Of course, the suspect might claim that they re-image their machine regularly as a matter of course – and not in response to legal action. This is, however, highly atypical behaviour in most circumstances, and may be looked-on with a jaundiced eye by any court.

Theoretically, data from an entirely overwritten disk can be recovered, although there is some argument among computer scientists about this point.⁵ Unfortunately the technology needed is beyond the scope of most forensic service providers (the process typically needs a scanning magnetic force microscope) and is inevitably very costly. We will discuss the recovery of erased data in more detail in a later article.

- **“Unfortunate accidents”.** Finally, if a suspect can afford the loss, there is no substitute for a well timed “accident” if they want to plausibly (or semi-plausibly) get rid of an unwanted evidential device. It is a well-known saying in IT security that a five-pound sledgehammer is the best disk wiping tool known to man. Unfortunately some of the people we investigate also discover this fact and apply it to the evidence in their possession.

In the last year we have seen an evidential laptop “accidentally” dropped down a flight of (concrete) stairs prior to collection, another shaken vigorously while it was switched on

⁴ Indeed, the original “Evidence Eliminator” application could be complemented by another application “Evidence Eliminator Eliminator” which was claimed by its developer to remove the material the original program missed. Unfortunately for its users, not even this second program was entirely reliable.

⁵ Wright C. *Overwriting Hard Drive Data* <http://sansforensics.wordpress.com/2009/01/15/overwriting-hard-drive-data/> excerpted from Wright C., Kleiman, D. and Shyaam Sundhar R.S. *Overwriting Hard Drive Data: The Great Wiping Controversy* in the proceedings of ICISS2008.

(causing a catastrophic drive failure called a “head crash”) and a (large) cup of coffee spilled over a running computer causing the hard drive to short-circuit. In the first two cases the damage to the evidence was too severe to repair, in the third case we were able to recover the comprehensively incriminating evidence after a hard-drive rebuild.

Not even physical destruction of evidence can guarantee that it will not return to haunt the suspect. On one occasion the suspect cut up a number of evidential floppy disks with a pair of scissors, assuming that the data would be irrecoverable. In the end the data was brought back by sellotaping the individual pieces of the damaged disks to new intact floppies and reading the data normally. In this case the suspect was convicted and jailed.⁶

2 Counterfeiting

Another common method of covering up information on a computer from an investigator is to simply change it. It is, of course, perfectly straightforward to modify material stored on a computer, and printed documents show no trace of any modifications that might have been made. When documents are examined in their electronic state, however, recent changes are far more obvious.

Many standard file types by default contain large quantities of metadata⁷ which can provide an evidence trail of any modifications that have been made to the file in the recent past. For example, a Microsoft Word file contains a list of its previous saved names and locations, the last accessed, printed or modified times, as well as the number of edits the file has undergone, the total edit time and the name of the last editor. Analysing this metadata can often demonstrate that a file has been tampered with.

If a suspicious file is examined in the context of the computer on which it was last edited, even more evidence is potentially available. The operating system preserves a different parallel set of metadata against which the internal data can be verified. As discussed already, many applications including Microsoft Office make temporary copies of files under modification as a precaution in case the computer “hangs” or “crashes” during the editing process. These temporary files are often useful in tracing the drafting history of any document, and can clearly show up counterfeiting attempts. Because material can accumulate on the hard drive that can validate or put into question the evidential documents stored on computer systems, IT forensics and electronic discovery specialists greatly prefer copying the entire hard drive of an evidential computer rather than just the potentially relevant files.

Defragmenting software, already discussed, can also be used as a counterfeiting tool. As discussed already evidence eliminator or file shredder software leaves considerable trace evidence of their use on hard drives. An observant investigator will find large areas of blank or pseudo-randomised data (called “data voids”) on the disk which can lead to a finding of deliberate evidence destruction. A defragmenter run may move existing files around the empty spaces of the disk, potentially creating many deleted copies of existing files thereby “repopulating” the erased areas of the unallocated spaces. This will make it a lot more difficult for an investigator to recognise that evidence eliminating software has been used.

⁶ Karl J. Flusche “*Computer Crime and Analysis of Computer Evidence*” in Rebecca Herold (ed.) “*The Privacy Papers*” Auerbach 2001

⁷ Metadata, often called data describing data, “is structured information that describes, explains, locates, or otherwise makes it easier to retrieve, use, or manage an information resource” – “*Understanding Metadata*”, <http://www.niso.org/publications/press/UnderstandingMetadata.pdf> NISO Press 2004

Fortunately, using the defragmenter can be seen as evidence destruction in its own right,⁸ and an experienced forensic investigator should be able to demonstrate that it has been used.⁹

Users attempting to “counterfeit” empty space have other options. One is to make copies of large folders on the computer and then delete them. This will show up clearly on the computer if the investigator is looking for it. Another option is to visit large numbers of web-pages using the web-browser, which will fill the empty spaces with browser cache remnants, or to download large media files from the Internet. The disadvantage with these approaches is that if the counterfeiter is not normally a heavy user of the web, or is not a regular downloader, this behaviour will appear anomalous to an investigator and will encourage them to pay particular attention to it.

Evidence counterfeiting does not just end with the use of word-processors and web-browsers. In one case the person under investigation decided to swap their entire hard drive with a new unit, set the clock on their computer back three years, install the operating system, then set the computer forward in time again. This gave the impression of a computer set up three years in the past and not used thereafter. The subject claimed that the computer has barely been used, and there was no data of relevance on it. The subject, however, neglected to observe that the manufacturing date of the replacement hard drive was two years later than their fictional “installation” date, which shattered the credibility of their testimony.

3 Data hiding

Hiding data is perhaps the least obvious counter-forensic strategy, but is more common than is widely appreciated. It relies on the belief that investigators typically have a limited amount of time or resources to devote to a single investigation, and an enormous amount of material to search through.

Modern computer hard drives are huge, capable of storing the data equivalents of million of printed books¹⁰ and so it is entirely practical to hide data on a modern computer or network. If the location or nature of a file is sufficiently obscured that commonly used data searching techniques will not detect it then data has been successfully protected against discovery by an investigator.

The disadvantage of data hiding is that the hidden data still remains on the evidential computer or network and a skilled, diligent or lucky investigator may still bring it to light.

Among the approaches we have seen in the past have been:

⁸ Lange M.C.S. & Nimsinger K.S. *Electronic Evidence and Discovery: What Every Lawyer Should Know*, American Bar Association, 2004.

⁹ ...and, crucially, that the use of the defragmenter is exceptional on the evidential computer. A standard defence used by individuals using the defragmenter is that they use it habitually to improve the performance of the computer. The defence is given strength by the fact that Microsoft have published technical documents stating that regular defragmentation is good practice (it is, but in most cases it is **something that needs to be done on an annual- or semi-annual- basis rather than weekly.**) More modern Microsoft operating systems also run defragmentation software automatically as part of normal operations, so investigators need to be able to tell automatic and manually initiated defragmentation runs apart. (<http://support.microsoft.com/kb/942092>)

¹⁰ The printed collections of the US Library of Congress can in principle be stored half-a dozen high capacity hard drives, costing no more than a few hundred Euros total. The University of California at Berkeley estimate the size of the printed collections of the US Library of Congress to be around 10 Terabytes (10,000 Gigabytes). A 1 TB hard drive costs, at time of writing, around €65.

- **Relocation of data.** The simplest way of hiding data is to move it to a separate storage location that may not be examined by an investigator. The disadvantage to this approach is that computers often record instances of data transfer in automated logs, potentially tipping off an investigator that data has been moved. Computers also typically update their system logs whenever a data storage device is connected to a computer, another way in which an investigator may spot that something is amiss.
- **Modification of file extensions.** Different types of computer files can normally be identified by their (usually three letter) “file extension” at the end of their file name such as .xls for an excel file.

One of the oldest methods of hiding a file from scrutiny is to alter its file extension. For example, to hide a JPEG graphics file an individual might alter its extension from .jpg to .exe, making it appear to be a program file rather than a picture. Windows identifies files by their file extension and will not, for example, open a picture file unless its file extension identifies it as such.

Fortunately, most modern forensic applications allow investigators to “signature analyse” all the files on an evidential computer to make sure that their file extensions match their actual file types. Most forensic investigators will do this routinely as part of any investigation, and will pay close attention to files whose signature doesn’t match their file extensions.

The signature is often based on values of the first few bytes of a file (sometimes referred to as the “file header”) , and occasionally, the last few bytes as well. These are often standard to a particular file format e.g. the first 10 bytes of a jpeg file are always FF D8 FF 60 00 10 4A 46 49 46, and D0 CF 11 E0 are always the first 4 bytes of a Microsoft Word .DOC document.

- **Cloaking of data.** Data is compressed, encoded or encrypted so that it is not found using standard search approaches such as keyword searches. Cloaking is far from foolproof, however, because the more effectively data is cloaked the more obvious the fact of the cloaking becomes to an investigator.

For example, zipped files are very common on evidential computers, but provide very little protection for the data residing in them. Conversely, encrypted files are usually very difficult to open without their encryption keys and are good protection for the data stored within them. These are relatively infrequent on evidential computer, however, and when found diligent investigators will pay special attention to them.

- **Encapsulation of data.** The most effective way of hiding data is to render it “invisible” – to obscure the fact that the hidden data even exists. Encapsulation hides files inside other, larger, files, making the presence of the hidden file difficult or impossible to detect. The two most effective approaches involve Steganography and Streaming.

Steganography^{11,12} (from the Greek for “hidden writing”) is a topic that has had entire books devoted to it.¹³ In computer terms it involves replacing the redundant data in a file with the data you wish to hide.

¹¹ Berinato, S. *The Rise of Anti-Forensics* in *CSO Magazine*, June 8 2007
http://www.csoonline.com/article/221208/The_Rise_of_Anti_Forensics retrieved 2 December 2009.

For example, documents can be hidden in graphic files by making use of the fact that modern graphics formats can display many more colours than the human eye can actually discern.¹⁴ This extra colour information has no practical function and can therefore be invisibly replaced with other data. In practice, a Word file might be steganographically hidden “inside” a picture file without altering the file’s size or the picture’s appearance in any way simply by reusing the space normally occupied by the redundant data. Such hidden files are exceptionally difficult to spot. Steganographically hidden files are typically only detected when the user of the technique is careless enough to keep the program that performs the steganography on the evidential computer, tipping off the investigator to the fact that such files may be present.

- **Streaming** is less complicated.¹⁵ Some operating systems allow users to associate more than one file with a single file table entry. Computer scientists refer to this process as “streaming” files. (This term is most often associated with NTFS.)
- Most forensic tools will only display one file associated with a file table entry. Any other files are essentially “hidden” to the forensic investigator, lost in the unallocated spaces of the hard drive. They may still be spotted by keyword or other searches, but files can be compressed or encoded before streaming, rendering such techniques useless. Hackers often use streaming to hide their tools on systems they have broken into.

¹² Inch S., *A Simple Image Hiding Technique: What You May Be Missing* in *Journal of Digital Forensic Practice* Volume 2, Issue 2 April 2008, pages 83 - 94

¹³ Cole E., *Hiding in Plain Sight: Steganography and the Art of Covert Communication*, John Wiley, 2003

¹⁴ Each dot or pixel making up a picture on screen is made up of three colours of light, Red, Green and Blue. The intensity level of each of the colours in each pixel is controlled by the numbers stored in three bytes of data in the computer’s memory. Each number controls the intensity of one of the pixel Red, Green and Blue colour components. Every pixel in a computer image is associated with 3 different bytes of data. Each of the three bytes can have a value between 0 (No colour) and 255 (intense colour)

A dim red may have the byte settings 64 Red, 0 Green and 0 Blue – the red component is set to emit light at only 25% of its maximum intensity - while the green and blue components are switched off. A medium yellow might be 100,100,0 (red and green light mixed produces yellow light). Bytes set to the values 0,0,252 would cause the corresponding pixel to show a bright blue.

This method for controlling pixel colours, with each of red, green and blue having 255 possible intensity levels, allows any pixel to display over 15 million different colours.

However the human eye cannot differentiate between more than a few tens of thousands of colours. Evolution simply didn’t produce an organ capable of greater colour sensitivity. This means that we simply cannot see the difference between a 64,0,0 red pixel and a 65,0,0 red pixel or a 100,100,0 yellow pixel and a 101,101,0 yellow pixel. The human eye is particularly insensitive to shades of blue so it probably cannot see the difference between a 0,0,252 blue pixel, a 0,0,253 blue pixel or even a 0,0,255 blue pixel.

Written in binary numbers 64 is 01000000. 65 is 01000001. Similarly in binary 100 is written 01100100. 101 is 01100101. You can see that in the case of red and yellow bytes it doesn’t matter whether the right-most or “least significant bit” in each byte is a 0 or a 1. The human eye cannot tell the difference in color. In binary, 252 is 11111100, 253 is 11111101, 254 is 11111110, and 255 is 11111111. Therefore, in the case of blue it doesn’t matter whether the values of the right-most two bits are 0 or 1. The eye cannot tell. This data is redundant.

This means that each pixel in an image has 4 bits of redundant data that can be used for something else, one in the red byte, one in the green and two in the blue. Two pixels between them have 8 redundant bits – a redundant byte. An image may be made up of thousands or millions of pixels, and so may contain kilobytes or megabytes of redundant data. This is more than enough to hide something else – invisibly. If you were to change this data for, say, a Microsoft Word file, the picture file would stay the same size, and the picture itself would show no discernable change, yet the Word file would be present and completely recoverable.

Some sound and movie file formats may contain redundant data in a similar way, and movie files can be gigabytes in size.

¹⁵ McClure S., Scambray J., Kurtz G. *Hacking Exposed (5th ed)*, McGraw Hill, 2005

4 Pre-emption

This technique involves stopping the evidence accumulating on a computer or network in the first place. In many cases it is difficult to stop computers from accumulating forensically useful material as would disrupt the normal operation of the computer.

There are some pre-emptive methods that can be quickly accessed and one of them is switching off the browser cache functionality on the Internet browser software. In most cases switching off this functionality does not stop the computer downloading or storing the data, however it ensures data is deleted when the browser leaves the web-site or is shut down. This means that some forensically recoverable data is stored on the computer, but in a form that is considerably more fragmentary and ephemeral than might otherwise be the case.

Pre-emption techniques are commonly used by computer hackers. Some tools such as U3 pens – USB thumb drives with complete suites of office and technical tools - are designed to prevent any data accumulating on a hard drive. Another approach is to run the computer using an operating system loaded from a bootable CD or DVD instead of the operating system on the hard drive. These “CD Distros” as they are known prevent anything being written onto a hard drive.

The problem with this pre-emption is that it is atypical behaviour. Ordinary users normally don't reorganise their disks to minimise slack data, or switch off the browser cache, usually because it makes the computer perform less efficiently. Similarly, most computers collect logging information by default, and even the largest logs fill only the tiniest fraction of the hard drives total storage capacity. There is no practical reason to turn them off other than to hide your activity on the computer.

Hence the problem! Pre-emption is inherently suspicious behaviour. It will tell the investigators that something “fishy” is going on and will encourage them to look at the computer in a great deal more detail.

Conclusion

Counter-forensic tools have become both more common and easier to use.¹⁶ Counter forensics has for years been a problem in the field of hacking investigations, largely because the malefactors are usually highly IT literate and are well aware of the kind of trace evidence their hacking tools are likely to leave behind.

Although many counter-forensic techniques exists and in recent years most powerful and complex such tools have been written by hackers for hackers, none of these methods guarantee to eliminate all evidence of their use and they are more likely to leave some traces of their use on computers, networks and servers.

Every method used to destroy evidence, once found, entails a risk to the individual doing the destruction. Any attempt to destroy evidence indicates that there is something to hide, and is

¹⁶ Berinato, S. The Rise of Anti-Forensics *CSO Magazine*, June 8 2007
http://www.csoonline.com/article/221208/The_Rise_of_Anti_Forensics retrieved 2 December 2009.

likely to destroy the credibility of the side that does it. There are few pieces of evidence that can be found on a computer that are more damaging than a finding of evidence tampering.¹⁷

Contact

Andy Harbison

Director, Forensic & Investigation Services
D +353 (0)1 6805 766
T +353 (0)1 6805 805
E andrew.harbison@grantthornton.ie

24-26 City Quay, Dublin 2

Noemi Kuncik

IT Forensics Specialist
D +353 (0)1 6805 662
T +353 (0)1 6805 805
E noemi.kuncik@grantthornton.ie

Offices also in Limerick and Kildare

This briefing is provided for general information purposes only and is not a comprehensive or complete statement of the issues to which it relates. It should not be used as a substitute for advice on individual cases. Before acting or refraining from acting in particular circumstances, specialist advice should be obtained. No liability can be accepted by Grant Thornton for any loss occasioned to any person acting or refraining from acting as a result of any material in this briefing. Grant Thornton, Irish member of Grant Thornton International, is authorised by the Institute of Chartered Accountants in Ireland to carry on investment business. www.grantthornton.ie

¹⁷ That being said, particularly in criminal cases, the penalty for evidence tampering may not be as severe as being caught with the original evidence, so there will often be an incentive to tamper. See Kuncik N & Harbison A. *A Brief Introduction to Counter-Forensics*, *Digital Forensics Magazine*, Issue 1, September 2009