

The Impact of Hard Disk Firmware Steganography on Computer Forensics

Iain Sutherland

Faculty of Advanced Technology
University of Glamorgan
CF37 1DL
+44(0)1443 654085
isutherl@glam.ac.uk

Gareth Davies

Faculty of Advanced Technology
University of Glamorgan
CF37 1DL
+44(0)1443 654085
gddavies@glam.ac.uk

Nick Pringle

Faculty of Advanced Technology
University of Glamorgan
CF37 1DL
+44(0)1443 654085
npringle@glam.ac.uk

Andrew Blyth

Faculty of Advanced Technology
University of Glamorgan
CF37 1DL
+44(0)1443 654085
ajcblyth@glam.ac.uk

ABSTRACT

The hard disk drive is probably the predominant form of storage media and is a primary data source in a forensic investigation. The majority of available software tools and literature relating to the investigation of the structure and content contained within a hard disk drive concerns the extraction and analysis of evidence from the various file systems which can reside in the user accessible area of the disk. It is known that there are other areas of the hard disk drive which could be used to conceal information, such as the Host Protected Area and the Device Configuration Overlay. There are recommended

methods for the detection and forensic analysis of these areas using appropriate tools and techniques. However, there are additional areas of a disk that have currently been overlooked. The Service Area or Platter Resident Firmware Area is used to store code and control structures responsible for the functionality of the drive and for logging failing or failed sectors.

This paper provides an introduction into initial research into the investigation and identification of issues relating to the analysis of the Platter Resident Firmware Area. In particular, the possibility that the Platter Resident Firmware Area could be manipulated and exploited to facilitate a form of steganography, enabling information to be concealed by a user and potentially from a digital forensic investigator.

Keywords: Digital Forensics, Hard Disk Drive, Firmware, Steganography.

1. INTRODUCTION

One of the main forms of storage media and therefore sources of data for forensic analysis is the hard disk drive. Current forensics practice for the most part deals with the analysis of hard disk drives found in server, desktop and laptop systems. It is also now becoming common for these drives to be embedded in certain other devices such as CCTV systems, games consoles and entertainment systems. The majority of the literature available on the forensic investigation of these devices concerns the analysis and extraction of evidence from the numerous forms of file systems which may reside in the user accessible area of the disk. It is known that there are other manufacturer areas of the drive which could be used to conceal information. However there are further areas of the drive that could be manipulated to enable data to be concealed from a user and potentially from a forensic investigator.

The focus of this paper is to highlight the potential problems that could arise from the manipulation of disk firmware, enabling the possible concealment of information on hard disk drives. Currently this technology is available mainly in data recovery labs that have specialised hardware to deal with the problem of faulty hard disk drives. However the technology is becoming available to the end user and as such should be of concern to the forensic investigator.

2. THE HARD DISK

The most common form factors are the 3.5 inch and 2.5 inch disks found in desktop and laptop systems respectively, although smaller versions can be found in other devices such as older versions of the Apple iPod. The disk may be housed either within the system unit or in a removable caddy. The main types are ATA and SCSI (Small Computer Systems Interface).

A hard disk drive is a complex device and can be viewed as a small computer system in itself. The hard disk drive is composed of platters, voice coils, read / write heads, casing, mountings, a motor, and a controller board. The data area

consists of a stack of metal, ceramic or glass platters coated with a magnetic film. One rotation of the disk at a particular radius is known as a track. For sets of surfaces, a set of tracks at the same radius is known as a cylinder. A separate armature and head assembly is present for each disk surface (one surface may be used to control the position of the read/write heads). The sector is the smallest addressable unit. A specific sector address can be found using the cylinder address (C) the Head (H) and the Sector (S). At a higher level of abstraction the Logical Block Address (LBA) method assigns a sequential address to each sector, but which may not relate to the block's physical location.

3. USER ADDRESSABLE SPACE

Hard disk maximum sizes are currently in the region of 1TB although speculations on higher capacities suggest the possibility of 4TB disk in the near future. The continuously increasing size of the hard disk drive poses a problem for the potential forensics examiner as it provides an ever increasing search space to examine as part of an investigation. Also with the increasing capacity and the decreasing cost of this form of storage device it is now not uncommon to find two or possibly more drives in a single home system, further increasing the problem for the forensic examiner. In the case of most drives, once the drive has been formatted and a file system is put in place the typical amount of storage is slightly less than the stated capacity. The forensic investigator has to be very careful to ensure they are aware of the correct storage capacity of the drive and have accessed all of the available information present on the storage media. This may include areas not addressable by the average user.

4. NON-USER ADDRESSABLE SPACE

Not all areas of the disk are addressable by the host computers operating system. In addition to the user addressable space there are areas of the drive that are used for the manufacturer to record data and perform diagnostics. These include the Host Protected Area (HPA) and Device Configuration Overlay (DCO), either or both of which can exist on a hard disk [1], [2] and in addition to this the firmware areas, which exist in the lower range of the address space, sometimes referenced by negative numbers (see Fig.1 below).



Figure 1. Overview of Disk Areas

4.1 Host Protected Area

The Host Protected Area (HPA) is used for holding diagnostics and other utilities required by the PC manufacturer (Gupta 2006). The presence of an HPA can be identified by commands `READ_NATIVE_MAX_ADDRESS` which provides the total number of sectors on the disk and `IDENTIFY_DEVICE` which provides total sectors a user can identify. Any difference between these two values indicates an HPA is present on the device [2].

4.2 Device Configuration Overlay

A Device Configuration Overlay (DCO) is similar to the HPA, but is used by manufacturers to configure drive sizes [3] and may exist at the same time. If a DCO is on the device this can be detected by the difference in values returned from the following two commands; `READ_NATIVE_MAX_ADDRESS` and `DEVICE_CONFIGURATION_IDENTIFY`. An excellent overview of HPA and DCO are provided in Carrier [2].

4.3 Firmware

In addition to the HPA and DCO there are other areas of the drive which are not addressable by the average user, but are vital to the disks correct operation. The firmware area is essential for correct drive operation. The firmware is composed of a series of modules, examples are; SECU (Security System Module), P-List, G-List, T-List, SMART Attributes, and U-List (Firmware Zone Translator). A portion of the disk firmware is resident on the drive platters, this is loaded by code located on the controller board of the hard-drive. Therefore these modules are located on either a) the hard-drive in a zone that is not normally accessible by the user or operating system, and b) into flash memory located/embedded on the disk controller. The firmware controls all aspects of the internal hard drive operation from system startup:

- On start-up, when a hard drive is powered on, the controller board loads the firmware modules from the disk platters into memory ready for execution. The firmware that is responsible for correctly configuring the hard drive and putting it in a ready state; providing all elements of the disk are working correctly the disk then presents itself as ready and allowing the host PC to load the Operating System.
- During the operation of the hard drive it is the firmware that ensures the correct operation of the hard drive, allowing it to correctly interact with other components on the system.
- When the hard drive is powered down, a shutdown sequence is executed by the firmware that ensures the hard drive powers down correctly so that it will operate successfully the next time it is powered on.

Various companies such as IBM, Hitachi and Western Digital provide software that can be used to up-date the firmware located on a hard-drive, sometimes allowing for more efficient operation.

5. FIRMWARE OPERATIONS

Firmware performs a number of key functions; one of these is SMART (Self-Monitoring, Analysis, and Reporting Technology) logs. As part of the ATA 3 standard there are a number of criteria which are monitored and logged as “threshold not exceeded” or “threshold exceeded” Attributes include read error, seek error, temperature, drive operation time etc. These Self-Monitoring, Analysis, and Reporting Technology (SMART) logs are aimed at predicting drive failure. The SMART attributes monitored depends on the manufacturer, for this reason they are usually of little forensic use to the examiner due to the differing implementations of the SMART log criteria in the different disks.

The firmware is also responsible for monitoring defect control, no disk is manufactured without some flaws and there will be some sectors on the drive which cannot be used. This process is transparently handled by the hard-disk and occurs ‘beneath’ the operating system level via the two lists, P and G [4]. Flaws identified on the drive during production are recorded in the disk firmware as the ‘P’ (permanent / primary / production) list. As the disk ages and through wear & tear other sectors may fail; this is recorded in the ‘G’ (growth) list. Reads and writes are automatically redirected (remapped) to spare sectors. P-list and G-list sectors are automatically bypassed by the drive electronics, and P/G-list sectors do not slow down drive access. By adding or removing a sector from the P-List and/or G-List we have the ability to hide/make-visible data on the hard-drive. Tools such as HDD Bad Sector Report [5] support a limited set of commands to modify/zero the P/G-lists.

6. TOOLS AND TECHNIQUES

To date the tools required to perform significant modifications of firmware are fairly expensive and found mostly in data recovery laboratories. However the authors are aware of two major commercial products available for this type of analysis and modification/repair of hard disk drives. Both sets of equipment were examined and comprise a combination of hardware and software tools. One coming from Russia and costing in the region of \$3000 for a UDMA set of tools. The full suite which includes the ability to extract data and work with some solid state devices and SCSI disks is in the region of \$15,000. A cheaper and more readily available device is offered from China and can be obtained via resellers in Europe for around \$300 per disk type (manufacturer). There are also a number of free/share tools that claim to read portions of the firmware, most commonly the disk serial number [6].

7. A SIMPLE STEGANOGRAPHY EXPERIMENT

This paper focuses on the possibility of concealing information on the disk via

the manipulation of the firmware and examines the issues and ease of use of some of the tools and techniques available. The authors are aware of a limited amount of material on this problem [7] although the potential for this particular form of drive behaviour was highlighted a number of years ago [8]. The ability to manipulate the disk firmware raises a number of potential issues which are demonstrated by the following experiment:

A 3.5" Fujitsu Hard Disk Drive (Model MPA3035AT) was forensically wiped, overwriting the contents with zeros. The disk was then reformatted with a partition and populated with a NTFS files system and a range of typical files; .doc, .txt, .jpps etc. Although the additional files are not directly related to the experiment they demonstrate the type of digital environment in which this type of steganography might be performed.

One of the text files (.txt) present on the drive was chosen at random and edited to include a distinctive keyword (a combination of the author names). This keyword, under normal circumstances residing within a document contained in the file system, would be easily located during a forensics analysis using the search facilities included in a number of commercial forensic software tools.

One of the firmware recovery tools outlined above was used to view the drive contents and also to locate the physical sector in which the selected text file was located. The particular model of the Fujitsu disk selected for this experiment supports two error lists in the firmware; one firmware list relating to production defects and another list relating to failing tracks on the drive. Any modification to the production list appears to require a reformatting of the drive. The P-list was not used in this experiment as our goal was to attempt to prove this form of steganography can be achieved on a 'live' system. Therefore the firmware error list relating to defective tracks (T-list) was modified to include an additional entry relating to the physical location of the modified text file.

The disk was rebooted and mounted. The firmware modification tool was then used to view the drive contents in an attempt to access the target file. The drive could no longer access the physical location (hidden data area) nor the data residing at that location. This was also confirmed via external hex editors, what is more the keywords were not present in any searches performed on the drive. The data is inaccessible by the disk drive and the computer operating system.

The firmware recovery tool was used to edit the error list returning it to its original state removing the previously added entry. The data area and text file containing the keyword was accessible on the drive.

8. DISCUSSION ON FORENSIC IMPACT

The findings of this initial experiment suggest that the use of this technique would permit a form of steganography enabling data to be concealed on the disk. The amount of space referenced by firmware error lists could be

substantial and span thousands of sectors storing a significant amount of information. This is in effect a more sophisticated version of marking a sector as bad on a floppy disk, as the operating system cannot access this portion of the disk.

This has the potential for a significant impact on computer forensic practice. It appears that tools commercially available for \$300 can be used to conceal information from a forensic investigator in a process that is relatively easy to accomplish, but difficult to detect. Standard forensic software cannot access the areas of the disk marked as 'bad' by a disk firmware error list. The manipulation of firmware is also not limited to the error lists [4]. In addition to the areas highlighted above disk firmware also contains the instructions for performing the LBA=CHS mapping, converting the LBA to the actual CHS locations on the disk. Standard forensic tools such as 'Encase' and 'AccessData' rely on the firmware translator operating correctly. They simply read all of the LBA's provided by the translator and create a corresponding disk image. The impact for forensics is that if the mapping has been tampered with or altered in some way then the data will not be present in the forensic image.

These problems suggest that in certain cases a standard forensic image may no longer be sufficient to fully analyse a hard disk drive and it may be essential to perform additional analysis on the original media. The analysis of firmware could be problematic for a number of reasons compounded by the fact that manufacturers implement the firmware in different ways on different models. This then poses a problem in the analysis of a disk if the validity of the firmware locations is to be assessed. The investigator will require a firmware tool to attempt a number of options:

Obtaining an identical model of disk will permit the validity of certain areas of the firmware on the suspect's disk to be checked. This can also be accomplished by comparing the suspect disk to a database of known firmware. This will detect the modification of some portions of the firmware such as the code that translates an LBA to a physical disk location [4]. However, the usual forensic practice of hash comparison would necessitate an MD5 hash for each and every firmware component.

Some of these components, in particular the error lists the target of this form of steganography are unique to each drive from the point of manufacture and evolving as the disk wears during normal operation. It is therefore impossible to validate the error list by comparison to another drive. Removal of the error list also presents problems as this may render the drive unusable.

The potential impact of the malicious functionality of the firmware both as a result of correct or malicious operation can also impact on information security. When wiping the disk the sectors in the error lists are not seen by the operating system so data may be left on these bad sectors. Disk disposal has

been documented as a significant issue [9], [10], [11] and traditional disk disposal tools may not be aware of this feature of a disk drive. The firmware in these disks may fail. The G-list may become full on some disk models and as a result the disk may stop working. An error in the firmware can prevent the disk being accessed while still physically healthy. Users with access to the appropriate tools and technology can now recover the data with relative ease. These users may be forensic investigators or those with less honourable intentions; potentially commercial competitors, foreign states engaged in commercial espionage although the technology is now becoming available to private individuals.

9. FUTURE ISSUES FOR CONSIDERATION

It is probable that the tools used to manipulate firmware and accomplish this form of steganography will become more readily available. If this is the case then future work should be focused in a number of key areas: It is suggested that work should also be undertaken to determine forensic best practice and procedure in this area. In some cases it may not be sufficient to work on an image of the disk drive. In particular it is of concern that this may be exploited as a possible route to introduce malware into a computer system. There are also a number of potential avenues that should be explored for solutions including the standardisation of some aspects in the way which disk firmware is structured or the generation of a library of trusted and true firmware which a forensic investigator can use to compare a suspects disk, although this would need to be a substantial library to cover all of the available models of hard disk drives.

10. SUMMARY AND CONCLUSIONS

This paper has discussed the implications of the malicious modification of firmware. It has suggested this raises a number of issues in the area of forensics and information security. In particular this route for steganography does not appear to be detectable by current forensic tools. In terms of information security working disks should make use of the ATA Secure erase function to ensure the secure removal of information from a disk. Disks that are faulty and which fail to be recognised by the operating system may not be truly 'dead'. This suggests disks considered dead or faulty should be put through some form of physical disposal system

ACKNOWLEDGEMENT

The authors would like to thank the members of the Information Security Research Group (IRSG) and in particular Mr. Konstantinos Xynos and Mr. Simon Harries.

AUTHOR BIOGRAPHIES

Dr. Sutherland is Reader of Computer Forensics at the Faculty of Advanced

Technology at the University of Glamorgan. His main field of interest is computer forensics, he maintains the University's Computing Forensics Laboratory. Dr. Sutherland has acted as an investigator and consultant on both criminal and civil cases. In addition to being actively involved in research in this area and supervising a number of Ph.D. students, Dr. Sutherland teaches Computer Forensics at both undergraduate and postgraduate level on the university's computer forensics degree schemes.

Mr. Gareth D. O. Davies is a Ph.D. Student at the Faculty of Advanced Technology in the University of Glamorgan. The main focus of his research is the security and forensic analysis of hard disk technology. He is a part-time lecturer on the Computer Forensics undergraduate degree at Glamorgan University and has been involved in a variety of other research projects in the area of Computer Forensics. Mr Davies has also acted as a consultant and assistant investigator on disk recovery technology cases in the University's Computing Forensics Laboratory.

Mr. Nick Pringle is a part-time Ph.D. Student at the Faculty of Advanced Technology in the University of Glamorgan. The main focus of his research is forensic analysis of large data sets. He has been involved in a variety of other research projects in the area of Computer Forensics, notably hard disk recovery. Mr Pringle has also acted as a consultant and assistant investigator on disk recovery technology cases in the University's Computing Forensics Laboratory.

Professor Andrew Blyth is Head of the Information Security Research Group at Faculty of Advanced Technology. His main interests lie in the area of Computer Network Management and Computer Network Defence and Computer Forensics. He has acted as an Expert Witness for National Police and Government.

REFERENCES

- [1] Vidström A., (2005) *Computer Forensics and the ATA Interface*, Technical report Swedish Defence Research Agency, FOI-R--1638—SE, February 2005, 1650-1942
- [2] Carrier B, (2005) *Forensic File System Analysis*, Addison Wesley.
- [3] Gupta M.R., Hoeschele, M.D., Marcus K. Rogers M.K., (2006) *Hidden Disk Areas: HPA and DCO*. International Journal of Digital Evidence, Fall 2006, Volume 5, Issue 1
- [4] Blyth A.J.C., Sutherland I, Pringle N., (2008) *Tools and Techniques for Steganography and Data Insertion onto Computer Hard-Drives*, 8th Annual Program Manager's Anti-Tamper Workshop, Sponsored by US DoD Anti-Tamper Executive Agent SAF/AQL and Department of the Army, Redstone Arsenal, Huntsville, AL, USA.

- [5] Badtrk (ADM) Documentation (Accessed 11/3/09)
<http://docsrv.caldera.com:507/en/man/html.ADM/badtrk.ADM.html>
- [6] HDD Firmware Serial Number Source Code 1.01 Free Download (Accessed 11/3/09)
<http://www.softlow.com/windows/development-tools/debugging/shareware/hdd-firmware-serial-number-source-code.html>
- [7] Davies G. & Sutherland I. (2009), *Forensic Implications of the modification of Hard Disk Firmware*, Proceedings of the Fourth Research Student Workshop, University of Glamorgan, 12th March 2009.
- [8] Gutmann .p (1996) *Secure Deletion of Data from Magnetic and Solid-State Memory*. Proceedings of The Sixth USENIX Security Symposium, July 22–25, 1996, San Jose, California, USA
- [9] Jones A., Valli C., Sutherland I. (2006) *An Analysis of Information Remaining on Disks offered for sale on the second hand market*. Journal of Digital Security, Forensics & Law. Volume 1, Issue 3.
- [10] Jones A., Dardick G., Sutherland I, Valli C., (2009) *The 2007 Analysis of Information Remaining on Disks offered for sale on the second hand market*. Int. J. Liability and Scientific Enquiry. Vol.2 (1), pp.53–68
- [11] Sutherland I, & Mee V. (2006) *Data Disposal: How educated are your Schools?*, 6th European Conference on Information Warfare and Security, June 2006.

