

2010

Zombie Hard disks - Data from the Living Dead

Iain Sutherland
University of Glamorgan

Gareth Davies
University of Glamorgan

Andy Jones
Centre for Information & Security Systems Research, BTUnited Kingdom

Andrew J. C. Blyth
University of Glamorgan

DOI: [10.4225/75/57b2b48d40ce3](https://doi.org/10.4225/75/57b2b48d40ce3)

Originally published in the Proceedings of the 8th Australian Digital Forensics Conference, Edith Cowan University, Perth Western Australia,
November 30th 2010

This Conference Proceeding is posted at Research Online.

<http://ro.ecu.edu.au/adf/86>

Zombie Hard disks – Data from the Living Dead

Iain Sutherland^{1,3}, Gareth Davies¹, Andy Jones^{2,3,4}, Andrew J.C. Blyth¹

¹University of Glamorgan
Information Security Research Group
Pontypridd, Wales, United Kingdom
isutherl@glam.ac.uk, gddavies@glam.ac.uk

²Centre for Information & Security Systems Research, BT
United Kingdom
andrew.28.jones@bt.com

³secau - Security Research Centre
Edith Cowan University
Perth, Western Australia

⁴Khalifa University of Science,
Technology & Research (KUSTAR)
Sharjah, UAE

Abstract

There have been a number of studies conducted in relation to data remaining on disks purchased on the second hand market. A large number of these studies have indicated that a proportion of these disks contain a degree of residual data placed on the drive by the original owners. The Security Research Centre at BT has sponsored a residual data study over the last five years examining disks sourced around the globe, in the UK, USA, Germany France and Australia. In 2008 as part of a 5 year study, Glamorgan University in conjunction with Edith Cowan University in Australia, Longwood University in Virginia USA and the BT Security Research Centre completed the fourth annual disk study aimed at assessing the volume and nature of information that remains on computer hard disks offered for sale on the second hand market. One of the main findings of the study was the high proportion of disks that are sold in a non-functioning state. As in both previous and following years a percentage of the hard disks examined in the 2008 study failed the imaging process and were marked as faulty. This paper describes further analysis of a number of these faulty drives from the UK sample set of the 2008 study. This paper details the analysis of non-functioning disks supplied to the University of Glamorgan to determine the ease with which data can be recovered from these drives using specialist recovery tools. It discusses implications for both computer forensics and information security practices and procedures.

Keywords

Forensics, disk disposal, data recovery

INTRODUCTION

There are numerous studies examining the way in which data can be recovered from computer hard disks for which the correct disposal action has not been taken (Garfinkel (2003), Sutherland (2006). During the last three years the University of Glamorgan in conjunction with Edith Cowan University in Australia, Longwood University in the USA and BT, has participated in a study of the data that remained on second hand hard disks (Jones et al, 2005, 2006, 2007, 2008 and 2010). The disk study project had been conducted in order to obtain an understanding of the amounts and types of information that remained on disks that had been offered for sale on the second hand market through various channels including online auctions.

The methodology employed for the disk study involved the purchase of a sample set of disks, in the case of the UK set in 2008 this was in the order of 133 disks, obtained from computer fairs, computer auctions or via online auction sites. The disks were purchased individually or in small batches, discretely obtained by a number of buyers. These disks were acquired by the project sponsor British Telecommunications (BT) Limited. The disks were a mixture of the form factors (the 3.5 inch desktop or 2.5 inch laptop drives) although a few smaller versions commonly used in devices such as the Apple iPod were also present in the sample. The disks were then supplied 'blind' to the researchers responsible for the imaging and analysis with no indication of where they had been sourced and identified only by a manufacturers label and a unique sequential serial number.

In the disk study, the disk drives were forensically analysed with the objective of determining whether the disks had been effectively cleansed of data or whether they still contained information that was either visible or easily recoverable, and if there was sufficient information to identify an organisation or individual. The research to date has indicated that a large proportion of the disks examined that could be accessed, still contained information pertaining to a previous owner of the disk. The information recovered could be considered of a sensitive nature to the original organisation or individual. The analysis used standard forensic good practice procedures as defined in the UK Association of Chief Police Officers (ACPO) Good Practice Guide for Computer-Based Electronic Evidence⁴⁵. Analysis was performed on a forensic copy of the drive, with the original disk being securely stored. The analysis methods used standard commercial tools and depending on the disk recovered information either from a file system remaining on the drive or from unallocated space by methods such as data carving.

The 2008 research also highlighted that a number of the disks failed during the imaging process; of the 133 disks analysed at Glamorgan for the 2008 disk study a total of 57 disks (44.19%) of the disks were faulty in some way. The figure of 57 disks does not include those that successfully imaged, but with a high rate of errors, which would be replaced with zeros in the image. The figure of 57 disks represents the number that failed to complete the imaging process using either Encase or FTK Imager, the main commercial forensic tools in use in the University of Glamorgan Forensics Laboratory. Prior to discussing the analysis and recovery of some these disks and considering the implications to computer forensics we outline some of the key features of hard disk technology and possible failures.

HARD DISK TECHNOLOGY

The hard disk remains the most popular form of digital storage media used in computer systems. The disk drive may be housed either within the system unit, a removable caddy or an external housing connected by either Firewire or USB. Current conventional computer hard disks are electro-mechanical devices. In addition to the Printed Circuit Board (PCB) the drive contains a number of moving parts: one or more metal, ceramic or glass platters coated with a magnetic thin film oxide (the data storage area within the disk), voice coils, motor, mountings and an armature with a separate head assembly for each disk surface. A specific sector (the smallest addressable unit) can be addressed by using the cylinder address (C) the Head (H) and the Sector (S). These mechanical components interact with very small tolerances and consequently are susceptible to damage from sudden movements, for instance the type of shock that occurs when a disk is dropped, even a relatively short distance. Some of the latest generation of disks contain integrated sensors on the PCB that detect sudden movements and automatically park the read/write heads in the event of an accidental fall. Once the heads are in the parked position some disks are rated to withstand shocks of up to 350G (Maxtor DiamondMax 20). Disks are therefore surprising robust in some cases but fragile in other cases.

Work conducted by Schroeder & Gibson (2007) suggested that disk independent factors in the form of operating conditions, usage and environmental factors affect the rate of failure rather than component failures resulting from the use of different types of disks, a conclusion supported by Pinheiro (2007). Problems may result from either mechanical or chemical failure, common issues highlighted in table 1 below.

Failure Class	Example Failure	Description
Mechanical failures	Spindle jam	This may result due to wear and tear on a pre-existing manufacturing defect – Semi-defective spindle bearing.
	Head crash	Insufficient airflow to maintain an air bearing between the head and the disk platter. This can occur due to a high impact during operation or a power failure. This can commonly occur as the disk powers down and result in the head failing to part correctly
Physical damage	Impact/ Shock	Incorrect handling and storage can cause damage to the disk. Physical damage to the circuit board can also result from a harsh working environment (Schroeder 2005)
	Disk contamination	Disks operate using very fine tolerances. The disk head travels over the surface of the drive at a distance of 0.00015 of an inch. Opening a disk in normal conditions can result in a degree of contamination that will eventually result in drive failure.

⁴⁵ Association of Chief Police Officers Good Practice Guide for Computer-Based Electronic Evidence, http://www.7safe.com/electronic_evidence/ACPO_guidelines_computer_evidence.pdf

Chemical	Chemical	A disk operating in a humid atmosphere may experience some deterioration in the electrical connectors between the circuit board on the disk
Electrical	Surge/ Component failure	Damage to voice coil, spindle motor can occur as a result of power surges.
Firmware/ Service area	Failure	Corruption of the firmware area may result due to incorrect storage but can also result due to the age of the device. It may also occur during disk use as firmware commands/modules corrupt the device

Table 1: Disk Failure

It should also be noted that some of the faults discussed above may arise as a result of malicious user activity or could also be the result of wear and tear on the disks.

FAULTY DRIVES PURCHASED IN THE 2008 DISK STUDY

The disks that were purchased as part of the disk study and that arrived in a faulty state could be as a result of three possibilities, each of which has a potential impact as to whether data had been removed from the drive:

- 1) The disk was sold in a working state and damaged in transit as a result of some form of kinetic shock. In this case the disks may or may not have had data removed.
- 2) The disk was knowingly sold in a faulty state after suffering some form of mechanical or electrical failure. In this case it is likely the user would not have had the opportunity to remove data from the disk.
- 3) The disk failed during its last use by the previous owner and was sold unknowingly in a faulty state. In this case the disks may or may not have had data removed.

Although the third option is a possibility, based on the number of drives faulty in the disk studies and the number of reliable power on / power off cycles, it is considered more likely that either the disks were damaged in transit or failed prior to being offered for sale. Most hard disks in a powered down and parked state are able to withstand a degree of mechanical shock. Although it should be noted that a number of disks that were obtained for the study were inadequately packaged when shipped from their previous owner. We therefore suggest that that it is reasonable to conclude that the majority of disks are most likely to fall into the second category those disks that have been knowingly sold in a faulty or failed state. If this is the case then the owners are unlikely to have had the opportunity to remove their data from the disk. A seller may be of the opinion that if they are unable to access the contents of a drive then neither can anyone else.

HYPOTHESIS

Based on these assumptions we suggest that comparing the sample of faulty drives in the disk study to those of the working disks; overall the faulty drives are more likely to contain user data and that provided the drives can be returned to a working state this data can probably be accessed.

HARD DISK FAILURE

A disk was considered faulty if the drive failed to spin up normally and present itself as a working disk to FTK imager, therefore resulting in it being not possible to access the disks using normal forensic tools. The disks when powered up for analysis and connected to a write-blocker displayed a variety of symptoms varying from spinning up (based on slight vibrations and an audible check), in some a rhythmic clicking of the disk heads, others produced no activity suggesting some form of electrical failure.

In one example, a disk from the 2008 disk study (07GB053) was reported as making an electronic pinging noise, similar to early 'space invaders' games when powered on. A visual inspection of the exterior of the disk revealed a slight dent in the upper cover in the area where the disk head traversed the platters of the disk. The disk was opened and it was determined that the read/write head was located in the centre of the platter rather than in the 'parked' position. As a result the depression in the outer case had resulted in one of the read/write heads becoming wedged on the disk surface. When power was supplied to the disk and it attempted to carry out the

initial read / seek this resulted in the voice coil⁴⁶ causing the pinging sound. The disk platters were manually rotated and the head moved back to the parked position. The disk was then powered on and after a single 'ping' the normal start up 'seek' sequence was performed with the heads travelling across the surface of the platter as normal.

However, subsequent attempts to recover data from the disk failed as it did not present itself as ready to accept/process ATA commands. The conclusion was that either the delicate read/write head mechanism had been damaged in some way as a result of it being forced down onto the surface of the platter, or that the disk spindle /platters had been distorted, although there was no visible damage to either the platters or disk heads.

A laser pointer was used to check for distortion on the platters. The resulting elliptical path of the reflected laser beam suggested that the spindle had been bent out of place and that this disk could only be recovered by relocating the disk platters into a new chassis, assuming the platter themselves had not been damaged. In this study, due to the difficulty in obtaining matching donor parts and to provide an idea of the capability of a user with access to only the equipment, repairs type 1 and 2 (Table 2) were applied so no further work was carried out on this disk.

It was determined that the faulty disks examined in the study could be divided into four classes in terms of the possible disk repair and data recovery, based on the skills and technology required.

Type	Repair
1	A simple mechanical fix, requiring no specialist tools or donor parts. Typically this would be resetting the disk head / parking the disk head / PCB repair.
2	Recovery using specialist tools either to image a failing disk drive or to repair disk firmware.
3	A complex mechanical fix requiring the use of donor parts to replace or rebuild faulty mechanical components of the disk.
4	No repairs possible

Table 2: Disk Repair

FAILED DISK ANALYSIS

The disks that did not image in the 2008 study were separated into SCSI and IDE drives, with the recovery tools available, this study focussed on the IDE drives which provide a sample size of 32 hard disks. The 32 faulty IDE disks were examined and none of the drives were determined to be type 1 drives, rather 10 drives (31% of the IDE disks) were identified as being type 2 and 20 (63%) drives were identified as being type 3 drives, requiring donor parts to return them to an operational state. Only 2 (6%) disk drives were judged to be irreparable, the pinging disk (07GB053) discussed above and one other disk drive which had undergone a head-crash which had caused significant damage to the disk platters.

The 10 type 2 drives were further analysed to determine the degree of data, if possible, that could be extracted from the disk drives. A more detailed analysis indicated that 3 (9%) of the disks were either partially supported by the recovery equipment and so could not be analysed further. The 7 disks (22%) that disks that could be repaired displayed two main problems: Either the need for a specialised interface and jumper requirements (4 disks) or varying degrees of firmware corruption (3 disks).

Restored Disk Results

Once the drives had been restored to enable access to the drive contents it was possible to determine that 2 of the 7 (29% of the restored disks) were wiped and that the remaining 5 (71% of the restored disks) contained varying degrees of data. This is in comparison to the 2008 disk study results where 31% of the readable disks were wiped and 69% contained data.

The data recovered from the restored drives included music, password files, surfing histories and for one company; financial contracts, work proposals, client details, email and attachments. Overall the results are broadly similar to the analysis of working disks in the 2008 study, but the sample size of 7 restored disks is very

⁴⁶ In a modern hard disk, the actuator uses a device called a voice coil to move the read/write head arms back and forward across the surface of the platters. The voice coil is operated by using electromagnetic attraction and repulsion.

small. Therefore future work will focus on the restoration of drives classed as type 3 (requiring spare parts) to determine if this makes a significant difference to the number of drives containing data.

CONCLUSION

The implications of this work suggest that disks that are found to be faulty and which fail to spin up or cannot be recognised by the operating system may not be truly 'dead'. A significant number of these disks can be restored and users with access to the appropriate tools and technology can then recover the data with relative ease. Many forensic investigators already utilise data recovery tools to recover potential evidence, however there may be other users with less honourable intentions.

Although the majority of the specialist tools described in this paper can be acquired for between a few hundred and a few thousand pounds, in terms of the threat vectors it is more likely to be commercial competitors from a large organisation, or nation states engaged in commercial espionage and is less likely to be from private individuals due to the cost and time involved in some of the recovery techniques.

REFERENCES

- American Forces Press Service (2006), Current Service members Possibly Affected by VA Data Loss, 6 June 2006.
- BBC News (2005), Data dangers dog hard drive sales, BBC, 12 September 2005.
- Canadian Globe and Mail (1993), Disk Slipped Into Wrong Hands, Canadian Globe and Mail, 2nd August 1993.
- Cullen D. (2000), Paul McCartney account details leaked on second user PC, The Register, 9th February 2000.
- Fragkos G. et al (2006) *An empirical methodology derived from the analysis of information remaining on second hand hard disks* in Blyth A and Sutherland I., *WFDIA Proceedings of the first workshop in Digital Forensics and Incident Analysis*
- Garfinkel S.L, Shelat A, (2003), Remembrance of Data Passed: A Study of Disk Sanitization Practices. IEEE Security & Privacy, Vol. 1, No. 1, 2003.
- Gutmann, P. (1996), Secure Deletion of Data from Magnetic and Solid-State Memory, Sixth USENIX Security Symposium Proceedings, San Jose, California, July 22-25, 1996.
- Gutmann, P. (2001), Data Remanence in Semiconductor Devices, 10th USENIX Security Symposium, Washington, D.C., August 13-17, 2001.
- Jenkins, C. (2005), Govt data sent to auction. The Australian, 2nd August 2005.
- Johannes, R. (2006), The Demographics of Identity Fraud: Through education and vigilance, banks can prepare and protect those most vulnerable, Javelin Research, http://www.javelinstrategy.com/uploads/607.R_2006_IDF_Demographics.pdf, Aug 2006.
- Jones, A., Mee, V., Meyler, C., and Gooch, J,(2005), Analysis of Data Recovered From Computer Disks released for sale by organisations, Journal of Information Warfare, (2005) 4 (2), 45-53.
- Jones A., Valli C., Sutherland I., Thomas P. (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.
- Jones A, Dardick G., Sutherland I, Valli C., (2007) *The 2007 Analysis of Information Remaining on Disks offered for sale on the second hand market*. Journal of Digital Security, Forensics & Law. IN PRESS
- Kerber R (2006), Firm will settle with state over data loss: Missing laptop had information on thousands, Boston Globe, 12 December 2006.
- Leyden, J. (2004), Oops! Firm accidentally eBay's customer database, The Register, 7 June 2004.

Price Waterhouse Cooper (2006), DTI Information security breaches survey 2006, http://www.dti.gov.uk/industries/information_security Sept 2006.

Pinheiro E., Weber W., Barroso L.A., (2007) *Failure Trends in a Large Disk Drive Population*, USENIX Association Storage Technologies, FAST '07: 5th USENIX Conference on File and USENIX Association Storage Technologies

Schroeder B., Gibson G. (2007) *Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?*, USENIX Association Storage Technologies, FAST '07: 5th USENIX Conference on File and USENIX Association Storage Technologies

Sutherland I, and Mee V. (2006) *Data Disposal: How educated are your Schools?*, 6th European Conference on Information Warfare and Security, June 2006.

Synovate, (2003), Federal Trade Commission – Identity Theft Survey Report, Federal Trade Commission, June 2006.

TechWeb, (2005), Seven-In-Ten Second-hand Hard Drives Still Have Data, TechWeb News, 31 May 2005.

Valli, C. (2004), Throwing out the Enterprise with the Hard Disk, In 2nd Australian Computer, Information and Network Forensics Conference, We-BCentre.COM, Fremantle Western Australia.

Valli C. & Woodward A., (2007) *Oops they did it again: The 2007 Australian study of remnant data contained on 2nd hand hard disk* Presented at the 5th Australian Digital Forensics Conference, Edith Cowan University Australia.

Vance A (2006a), Ernst & Young fails to disclose high-profile data loss: Sun CEO's social security number exposed, The Register, 25 February 2006.

Vance A (2006b), Wells Fargo fesses up to data loss: Lightning strikes twice for HP man, The Register, 12 May 2006.

Vidström A, (2005) *Computer forensics and the ATA interface*, White Paper, Swedish Defence Research Agency

ACKNOWLEDGEMENTS

The authors would like to acknowledge the generous sponsorship of British Telecommunications (BT) who support the disk study each year. In addition to the authors of this paper, we would like to acknowledge the other researchers who participated in all of the disk study work in current and previous disk studies and our colleagues and friends at Edith Cowan, Longwood and Glamorgan Universities.