

Solid State Forensics:
Investigating the Effects of Garbage Collection
on Potentially Volatile Data During the
Process of Forensic Extraction of SSDs

Author: Michael Lawson - C1542550

Supervisor: Michael Daley

Moderator: Kirill Sidorov



BSc Computer Science with Security and Forensics

One Semester Project - 40 Credits

Module Code: CM3203

2018

Table of Contents

- 1. Abstract 4
- 2. Introduction..... 5
 - 2.1 Intended Audience and Beneficiaries 6
 - 2.2 Project Scope..... 6
 - 2.3 Project Aims and Objectives..... 7
- 3. Background..... 8
 - 3.1 Brief History of Computer Storage..... 8
 - 3.2 The Features, Functions and Differences Between HDDs and SSDs 10
 - 3.2.1 Hard Drive Disks 10
 - 3.2.2 Solid State Drives 11
 - 3.4 Timeline of Significant Milestones in SSD History 12
 - 3.5 Forensic Recovery 13
 - 3.5.1 SSD controller 14
 - 3.5.2 Writeblockers 15
 - 3.5.3 Garbage collection and TRIM function 15
 - 3.5.4 Project Discussions with Gwent Police 19
 - 3.5.5 ACPO Good Practice Guide for Digital Evidence 19
 - 3.5.6 Recent Studies 20
 - 3.5.7 Conclusion for Background..... 22
- 4. Specification & Design 23
 - 4.1 Approach..... 23
 - 4.1.1 Details of Experimental Equipment..... 23
 - 4.1.2 Manufacturing Information and Details of Each SSD..... 23
 - 4.1.3 Forensic Bridge/Writeblocker 26
 - 4.1.4 Forensic Toolkit..... 26
- 5. Implementation 27
 - 5.1 Constraints 27
 - 5.2 Experiment Methods 27
 - 5.2.1 Creating the Image 28
 - 5.2.2 Experiment Preparation 28
- 6. Results and Evaluation 30
 - 6.1 Experiment Results 30

6.1.1 Experiment 1.....	30
6.1.2 Experiment 2.....	38
6.1.3 Experiment 3.....	50
6.1.4 Challenges and Alterations Faced During Experimentation	59
7. Conclusion	61
8. Future Work.....	64
9. Reflection	65
10. Glossary	66
11. References	67

List of Figures

Figure 1: Impression of the TRIM function processes (Gubanov and Afonin, 2012).....	17
Figure 2: Impression of the garbage collection function (Thatcher, 2009).	18
Figure 3: Screenshot showing that the TRIM is enabled	29
Figure 4: Kingston SSD Matching MD5 hash.....	31
Figure 5: Kingston SSD- files with zeros as data.....	32
Figure 6:Kingston SSD- Shows picture changes when recovered.	33
Figure 7:WD Green Time taken and matching MD5.....	34
Figure 8:WD Green Files been filled with zeros.....	35
Figure 9: Crucial SSD Time taken and MD5.....	36
Figure 10:Crucial SSD files been filled with zeros	37
Figure 11: Problems with experiment 2.....	40
Figure 12:WD Green- Files when first imported to FTK.....	41
Figure 13: WD Green-File displayed in natural form	42
Figure 14:WD Green SSD- Files after being processed	42
Figure 15: WD Green SSD- Side by side comparison of processed files	43
Figure 16: Kingston SSD-After FTK was paused	44
Figure 17:Kingston SSD- before FTK process.	44
Figure 18:Kingston SSD- After FTK processed files	45
Figure 19: Kingston SSD- Side by side comparison	46
Figure 20: Crucial SSD- before images processed in FTK	47
Figure 21:Crucial SSD- example of file in Hex format	48
Figure 22: Crucial SSD- Example of file in natural format	48
Figure 23: Crucial SSD- Example of file change after processing in FTK	49
Figure 24: Diskpart in CMD	51
Figure 25: Kingston SSD- files when loaded in FTK	52
Figure 26:Kingston SSD- Files after being processed by FTK	53
Figure 27: Kingston SSD- example of file with zeros.....	54
Figure 28 :Kingston SSD- File with its natural display.	54
Figure 29: WD Green SSD- Files when first loaded into FTK.....	55
Figure 30: WD Green- Files after they were processed by FTK	56
Figure 31:Crucial SSD-Files when loaded into FTK.....	57
Figure 32: Crucial SSD- Files after being proceed by FTK.....	58

List of Tables

Table 1: Results from experiment 1.....	37
Table 2: Expected results for experiment 2	41
Table 3: Results of Experiment 2	49
Table 4: Expected results of experiment 3	51
Table 5: Actual results of experiment 3	58

Acknowledgements

I would like to thank my supervisor Mike Daley for being available and seeing me when I had issues, even when it was short notice. I would like to take this chance to say thanks to my Mum for giving me continuous support throughout my ever-changing paths in life. Finally, "I" would like to thank Alice Sims for her support and patience throughout this dissertation

1. Abstract

As solid state drives increase in popularity throughout the entire computing world, the rapid technological advancements made to these devices have left the digital forensics community at odds with the manufacturers. Forensic recovery of traditional hard drive disks was a well practiced exercise as the methods of data recovery remained unchanged for many years. However, with SSDs boasting space and time saving functions, such as TRIM and garbage collection, this has altered the way in which digital forensics is performed for the foreseeable future. These functions are not fully understood and not even fully disclosed by the manufacturers and are the main causes of SSD self-corrosion. This has caused huge data loss issues for the digital forensics community and police forces that intend on gathering evidence from a drive. Therefore, this project has been written to provide further information, through research and experimentation, into the best ways in which digital forensic investigators can image SSDs as efficiently as possible. It was found that the presence of a write blocker helped stabilise the image, when connected quick enough and disabling the automount resulted in the recovery of more data.

2. Introduction

In the last decade, the world of computer memory storage has experienced significant advancements. The once, almost exclusively, used Hard Drive Disks (HDD) has been slowly replaced by the new innovative Solid-State Disks (SSD). The popularity of SSDs has risen due to their vast performance advantages over HDDs. They transfer data faster, make less noise, are more reliable due to lack of mechanical parts and are generally more compact, all of which has made them far more desirable for the everyday user. However, SSDs have proven that they pose notable issues for digital forensic investigators during data recovery. This is attributed by how SSDs process deleted data and how it is vastly different to HDDs.

When data is deleted from an HDD it is usually preserved until new data is overwritten to the same location, replacing and deleting the old data. Forensic investigators are generally able to recover parts of the data even when overwritten, this is called recovery from slack space. The process of creating slack space occurs when the new data is smaller than the older data which is to be overwritten. Often this data is available to recover. If a user wanted to prevent recovery or just prolong the life span of the HDD, software can be used to overwrite the deleted locations with zeros, leaving little for investigators to find.

In contrast, SSDs write data to multiple cells which together form a block, blocks are then grouped together to form a page. In SSDs it is not possible to edit or delete single cells, instead they perform wear-levelling. This is when data is automatically transferred from a block which has been well-used, to an under-used block. This prevents wear on cells which prolongs life span and ensures the space is well utilised. A TRIM command is then sent from the operating system which signals that the page is no longer being used and is subsequently deleted to create extra space (Gubanov and Afonin, 2016).

The TRIM command adds the original file location to a feature on the SSD called Garbage Collection. This functions without assistance from the Operating system and begins to find all pages associated to the file and deletes them in the background. Although this function is helpful for the user as it slows wear and creates additional space, it has caused issues for forensic investigators whilst attempting data recovery (Gubanov and Afonin, 2016).

In digital forensics one of the main intentions is to ensure integrity is maintained throughout the entire investigation. A duplication of original drive is created through a process called imaging. Once this is complete a mathematical hash is performed before the image is examined and after to guarantee no alterations have been made to the image throughout the investigation. An example of which is an MD5 hash, which is a mathematical function applied against a file to produce a unique value. If a file has been deleted, altered, or moved, it will produce a different MD5 hash. For this reason, it is important to MD5 hash the experiments conducted within this project.

HDDs use a write blocker to ensure nothing is written to the drive which would change the MD5 hash. However, with SSDs the Garbage Collection function runs independently of the operating system. Therefore, it is possible for it to run in the background and consequently alter the MD5 hash, further challenging the aims of a forensic investigator.

The challenges faced by forensic investigators working with SSDs have motivated the completion of this project. The question posed in this project was originally proposed by the Gwent Police who wanted to gain more insight into how these issues can be solved after experiencing difficulties when imaging SSD drives.

The concept for this project was initially suggested by Gwent Police, however, it has the capacity to benefit the wider digital forensics community. This is because the computer storage market is rapidly moving towards using SSDs, however, the more efficient the SSDs are becoming, the more difficult forensic recovery of SSDs is. Therefore, this project aims to contribute towards digital forensics of the future.

To fulfil the project outcomes the University of Cardiff provided all equipment necessary to conduct the experiments. The experiment was based upon testing 3 randomly selected makes of SSD, assessing the impacts of the Garbage Collection and TRIM functions to determine if forensic recovery of SSDs can be improved through minimising the effects of these functions, whilst maintaining forensic integrity throughout. The experiments will be performed using the popular data recovery program FDK, and will be carried out through real world testing, as it was anticipated that this would provide the most valid and reliable outcomes.

When interpreting the outcomes of this project, it must be considered, that as a university project there are limitations on the amount of experimentation that can be conducted without a team and there is a restricted amount of equipment available. Nevertheless, this project hopes to provide a basis for further research into the challenges faced by forensic investigators when recovering data from SSDs.

2.1 Intended Audience and Beneficiaries

This project was predominantly conducted to answer questions that Gwent Police had posed regarding the effects of TRIM function and garbage collection whilst investigating SSDs. Further reaching, it is hoped that this paper will also provide useful information for the digital forensic community and help other police forces and digital forensic professionals to understand how to effectively investigate SSDs.

2.2 Project Scope

The main issue raised by this project is to prevent SSD functions such as garbage collection and TRIM from permanently deleting data, making it extremely difficult to recover valuable data during a digital forensic investigation. This problem is relatively new and little research has been conducted to solve the issues, let alone understand the workings of these functions.

After researching the subjects in depth, the project outcomes have altered since the initial plan was submitted. Since there are multiple SSD manufacturers, all with varying ways in which they implement the garbage collection and TRIM functions, solving this problem entirely is not likely to be achieved within one project. It is also an ever-developing field of

computing, so the initial plan outcomes of finding and solving the cause of data loss was too ambitious given the time restraint within this project.

Therefore, the real-world experiments and research presented within this paper aim to grasp an understanding of how these functions are implemented and how they affect data loss, which will inform future studies. In addition, it has also contributed towards answering the main queries posed by Gwent Police regarding this subject.

2.3 Project Aims and Objectives

Primary

1. Determine the differences between the effects of each SSD's garbage collection and trimming processes on an image.

- This will be achieved by deleting files from each SSD and then attempting to recover them at different intervals. Once recovered, the differences between the recovered images will be compared against each other to see which ones performed best.

2. Understand what triggers garbage collection.

- Research will be conducted to gain insight on the technical processes of garbage collection to build an idea of the potential triggers.

3. Suggest ways of preventing garbage collection.

- After research and experimental practice, I will make suggestions for potential ways in which garbage collection can be prevented or controlled during forensic analysis.

Secondary

1. Suggest safe methods of imaging.

- The secondary aim has changed since the initial plan to suggest ways in which Gwent Police can image SSDs in a more effective way. This aim is more feasible after researching the questions posed by Gwent Police and seems to be a more useful research outcome for them.

3. Background

3.1 Brief History of Computer Storage

Since the invention of computers, people have continually developed ways in which to store data efficiently. Before the widespread usage of computers, the first attempt of storing machine language data was through the use of punch cards. The holes punched in the card represented a unique set of instructions which could be translated into a meaningful action such as automated pianos and textile looms. These cards were used in a variety of creative ways for two centuries, until the 1960s saw the gradual replacement of punch cards with 'magnetic storage' as the main method of data storage (Foote, 2017).

First developed in the 1920s for the purpose of recording sounds, magnetic tape revolutionised the way in which computer data was recorded, stored and used. It consists of a thin, plastic film that is magnetised along one side. The tape works by passing over tape heads from one reel to another, using motors, and information is written, erased or read as it moves (Stevens, 1981).

The computer manufacturing company, IBM, produced an innovative, high-speed magnetic tape system with a vacuum column design called IBM Model 726. This was released for sale in 1953 alongside the first electronic digital computer, IBM Model 701. It had the capabilities of storing much larger amounts of data than any system previous and was a significant development in the computer revolution (Stevens, 1981).

Magnetic storage technology continued to develop throughout the 1960s and became ever popular with the increasing accessibility of computers. During this time magnetic tape was still the primary method of computer storage. However, the development of magnetic disk storage introduced the advancing concept of removable memory. IBM introduced Model 1311 in 1962 which allowed for removable disk pack and potentially unlimited offline storage, depending on how many disk packs were purchased. The removable disk packs were designed to be cheaper, more portable and more compact than the magnetic tape systems (IBM, 2018).

In the subsequent two decades, disk drives continued to evolve, becoming physically smaller, but with larger capacity and more efficient processing. Then in 1980 Seagate Technology developed the first Hard Disk Drive (HDD), the ST506. It had five times the capacity of a standard floppy disk but was able to fit in the space of a floppy disk drive. It was made from a strong, metallic structure, coated with magnetic material to store the data. The ST506 fast became a popular product, with large computing companies such as Apple Computer and IBM becoming customers (Goda and Kitsuregawa, 2012).

The 1980s saw significant developments in computer technology alongside the influx of computers into schools, workplaces and some homes. People began to view them as accessible devices, rather than the enormous machines used only by large companies and scientists. This was due to scientists developing miniature computing components such as the microchips, to condense a computer into an appropriate size for a house (Wilkes, 1980).

In 1984, Fujio Masuoka invented flash memory, which proved easier and faster to store digital information. Flash memory was the founding technology of the solid state drive (SSD), consisting of no moving parts, with all processes conducted electronically instead of mechanically. It was developed from EEPROM (Electrically Erasable Programmable Read Only Memory) which means that data can be electronically erased and reprogrammed from the drive.

Although SSDs were beginning to evolve in their own right, HDDs were still prevalent and were advancing technologically. A breakthrough was made in 1990 with the formation of IBM 9345 hard disk drive. Codenamed the 'Sawmill', it grew in popularity beyond the competitors due to significant increases in capacity. This was factored to the implementation of magnetoresistive heads which allowed the device to store bits more densely (Computerhistory.org, 2018).

The 1990s proved to be a fast moving decade for computer advancements, from the announcement of the World Wide Web; advancements in graphics and gaming within PCs; the development of the Google search engine; to the introduction of Wi-Fi in people's homes. The fast paced innovation required computer storage devices to improve their capabilities to match the needs of the growing computer technology market.

By 2000 the introduction of the USB Flash drive offered 4 times the capacity of common storage devices at the start of the 1990s. They consisted of a small case, containing flash memory with a USB interface. They had great versatility and were designed to store, back up and transfer data between various devices. Their design also ensured their longevity in comparison to optical discs and floppy disks as they were resistant to scratch damage and resilient against magnetic erasure. These factors resulted in the reduction of optical disc and floppy disk drives installed in PCs and laptops, in exchange for USB ports (computerhistory.org, 2018).

With the growing need for computing systems to have more storage with more power, the early 2000's saw the development of the first 3.5 inch hard drive to contain 1 Terabyte of capacity, named the Hitachi DeskStar 7K1000. With smoother operation of mechanical parts, the reduced heat and wear increased the reliability and overall performance of the HDD. Models with similar technology were soon released from other companies, all with capabilities to create more powerful and efficient computer storage solutions for the growing games, networking and video editing markets (Hgst.com, 2018).

In more recent times, major developments were made in Cloud storage, with Amazon Web Services launching Cloud storage platforms that allowed users to pay only for the capacity they required. This reduced the need for companies to create their own complex storage systems and the extra space and expertise needed to uphold it. Other companies soon adopted Cloud storage concept and created similar services that are commonly used today, such as Dropbox and Google Drive (Rajan, 2012).

In today's world of various and complex computing needs, multiple storage types are used depending on the user's requirements, with Cloud storage, USBs, HDDs and SSDs and all still relevant. Different storage types facilitate different types of data and since the rise in popularity of smart devices and social networking, the average person now uses and stores significantly more data than they did just a decade ago. Therefore, computer storage companies are racing more than ever to develop smaller, more powerful data storage solutions for the increasing demand.

3.2 The Features, Functions and Differences Between HDDs and SSDs

Until recent years HDDs have dominated the world of computer storage. However, with the advancements made to SSD storage, they are growing ever popular and even beginning to overshadow the HDDs in performance and availability. At a time where both are still prevalently used, this section will outline the notable features and functions of both storage facilities and will discuss the differences between the two, with particular focus on the progressive development and use of SSDs.

3.2.1 Hard Drive Disks

HDDs have been in use since the 1950s. The first models were relatively large and had the maximum capacity of a few megabytes. However, after considerable advancements, HDDs are now able to store terabytes of data within small, portable devices.

HDDs are commonly found within desktop computers. They are approximately the size of a paperback book and are encased within a mounting, enabling them to be attached to the drive bay of a computer. Inside the case the mechanical parts are found, consisting of a single platter or stacked platters which spin around the spindle, the faster the spin the better the performance. By changing the magnetic fields on the platters, using the read-write head, information can be written to and read from the drive. From the back of the HDD there is a cable to connect to the motherboard. Depending on what type of HDD it is, the cable is either a SATA or PATA.

The methods used by a HDD to store data is described as non-volatile storage. This means that unlike volatile storage, such as RAM, the HDD is able to store data when the power is off. Therefore, the data can be accessed after a computer restart without any data loss. Features such as this allow users to store data for long periods of time, or in the context of this project, allow forensic investigators to recover and gather historic evidence from a computer system.

The major benefits of HDDs are their ability to store large amounts of data relatively inexpensively. Their technology is tried and tested, so their performances are predictable and reliable. They are easily accessible with the two most common forms the 2.5" (for laptops) and the 3.5" (for desktop computers). The larger HDDs allow more space for additional platters, which provide more storage capacity.

3.2.2 Solid State Drives

SSDs have a long history, but have only recently started to expand into the market of computer storage in a significant way. The concept technology of the SSD originates in the 1950s, with the card capacity read-only store and magnetic core memory. Although these memory units were outlining the future for SSDs, they were overshadowed by the more affordable drum storage units, resulting in their usage ceasing.

The SSD that is recognised as today's prominent computer storage unit evolved relatively recently with the introduction of flash-based SSDs in the 1990s. Flash memory is a non-volatile memory chip that retains information without the need for constant power. The combination of flash memory and SSDs is attributed to the success of the flash based SSD.

The implementation of NAND flash within SSDs gave them significantly higher performance rates, which lowered the power requirement and offered a higher capacity than the leading HDDs. With the continual technological development since the 1990s, prices of flash-based SSDs have dropped dramatically. This affordability has increased their popularity within the commercial markets.

The one obvious advantage of SSDs is that they can perform at much faster rates as they do not have the constraint of moving parts. Apart from this there are many other advantages to SSDs over using HDDs including using less power, generating less noise and heat, increased reliability and lack of affect by magnetism. However, although SSD prices continue to fall, HDDs used for storing larger quantities of data still offer a better value for money.

Solid-state drives are named that specifically because they differ from HDDs in structure, containing no spinning disks or other moving parts. Rather than relying on read-write heads to translate information like HDDs, the information is saved to a pool of NAND flash. So instead of containing the stacked plates like HDDs, the SSDs contain a grid of electrical cells that are used to send and receive data. The grid is separated into smaller sections, with their purpose to store this data, these sections are known as pages. The pages are then collected together to form blocks.

Each SSD contains an embedded processor called a controller. Its main function is to bridge between the NAND memory mechanisms and the computers using firmware level code. It is an important component that contributes greatly to the SSDs performance levels as it is responsible for the implementation of many tasks including reading and writing cache, encryption, garbage collection, error detection and wear levelling.

Another of the controller's key functions is to map out bad memory cells and locate spare cells for new data to be written to. Therefore, each time the computer needs to read or write data, it communicates with the controller and the controller translates that to the SSD, so the function can be carried out.

When the SSD memory is close to capacity, the controller searches for blocks that are no longer in use, these are known as stale blocks. These usually contain old data that is awaiting deletion before new data is able to write to that location. This process as an entirety is called garbage collection and is usually conducted independently of any other processes and when

the computer is idle, it can also be triggered, is when a host computer deletes files. Another way in which the controller enables the garbage collection function, is to search for any pages that relate to the data that is being saved and marks them ready for garbage collection, this process is called the TRIM function.

When receiving data, SSDs are only able to write it to empty pages in a block. In contrast, HDDs write data onto any location of the plate, meaning that data is easily overwritten. SSDs are unable to overwrite data directly. Understanding these internal processes of SSDs is highly important when addressing the main themes and questions of this project.

3.4 Timeline of Significant Milestones in SSD History

- The 1970s there existed rewritable, non-volatile, solid state storage modules, but it wasn't until much later, after numerous setbacks and false starts, that the SSD form known in today's computing was developed.
- 1977 saw the launch of MM-S100. A non-volatile RAM card with no moving parts. Although it lacked mechanical parts it wasn't described as an actual SSD as the memory was not implemented by semiconductors. However, it was an integral step towards the development of solid state memory systems.
- In the early 1980s, Intel created bubble memory. There was a lot of interest surrounding this new product, which was coined as a solid state floppy disk. The product failed to become cost effective and eventually it was discontinued.
- 1990 NEC began to market SSDs that used internal battery backed RAM.
- In 1991 SunDisk, which later became known as SanDisk, sold the world's first flash SSD to IBM. It was only 20MB but cost \$1000.
- Solid Data Systems was founded in 1993, who developed new technology that enhanced the performance of SSDs by creating Direct Addressing™. This function eliminated intermediate delays, meaning SSDs became faster and more efficient.
- By 1999, 11 companies had begun production and marketing of SSDs.
- The year 2000 saw the first online adverts for SSDs with various companies beginning to take advantage of the online advertising platform.
- In the May of 2001, Winchester Systems released a product called Flash SSD. By the October of this year, there were officially 21 manufacturers that were actively marketing SSDs.
- The first NAND flash SSD was released in 2002.
- In 2003 the first terabyte SSD became commercially available. In the August of this year, Ramtron built the first FRAM (ferroelectric random access memory). It is clear now that this year was a significant point in which SSDs could not be displaced within the modern market.
- In the May of 2005, Samsung introduces their first SSDs into the market and sets the trend for other large companies to follow suit.
- Samsung announce first prototypes of PRAM in 2006 (Phase-change Random Access Memory). Samsung predicted that it will replace NOR flash within a decade.

- By 2008, over 100 companies have begun their own manufacturing of SSDs, with competition increasing rapidly.
- In 2010 the first NAND flash memory is sampled by Intel and Micron. Companies also began to brand the SSD controller with SandForce Driven SSDs. This is good for companies that are attempting to creating unique versions of SSDs, however, secretive controller technology has contributed towards the forensic investigative issues surrounding SSDs since its development.
- In 2013, there are even more companies that creating SSDs. Therefore, higher rates of competition, which encouraged further secrecy of their technology
- By 2014 most large SSD manufacturers are designing and patenting their own controllers. This caused disputes as law suits between different companies when one would accuse another of copyright. A high-profile case was one of Netlist asking a US court to shut down SanDisk operations for copyright. Ultimately Netlist lost their case, but many other similar disputes occurred between other companies at this time.
- In the August of 2015 the first M.2 SSDs were developed using MRAM instead of flash.
- Due to the intense competition and race to develop better technology, 2016 saw a rise in some companies tricking consumers into thinking that they are buying better quality products than they were. By trialling their best products but swapping them out for lesser quality ones for the mass market. SSDs were also beginning to shrink to fit into mobile phones and tablets.
- Up to 2017, there has been a continuation of technological development across numerous companies and increased secrecy about the internal workings of their SSD functions.

3.5 Forensic Recovery

The craft of digital forensics was once an extremely specialised skill. However, in recent decades it has become increasingly important, used in various agencies, such as the military and police, to recover data and gather evidence. Although the world of forensic recovery has rapidly developed since its beginnings, so has the technology that is under forensic scrutiny. This has led to a variety of new issues that prevent efficient forensic recovery from occurring (Garfinkel, 2010).

The practice of digital forensics is approximately four decades old and was initially used to recover data and piece together files that had been lost, deleted or fragmented. This work was mainly conducted by computing experts working within law enforcement agencies, who performed these forensic recoveries only on rare occasions. This is because during the 1980s most disks could not hold any considerable amount of data, so perpetrators of crimes would more likely be caught by careless storage of information that had been printed out (Garfinkel, 2010).

Between the years of 1999 until the early 2000s were considered a “Golden Age” for digital forensics. This is because most computers were using HDDs to store large amounts of data. At this time, the workings of HDDs became common knowledge, so it was widely

acknowledged in the digital forensics industry that data deleted from HDDs appears to be erased, but in fact it remains in the same location until it is overwritten with new data.

This allowed the 'deleted' data open to recovery by forensic investigators. At this time, the internal mechanisms of data storage on HDDs was common knowledge and the techniques to extract such data was well practiced. However, the skills and knowledge gained during this time has become increasingly irrelevant as technology has progressed so dramatically and large companies developing SSDs now withhold information on the technicalities of their read write mechanisms to remain competitive. This rapid advancement and secrecy surrounding SSDs has hindered the process made by forensic investigators during the Golden Age of digital forensics (Garfinkel, 2010).

The introduction of computers containing SSD storage changed the predictable and usually reliable steps taken to recover deleted data and introduced more uncertainty to digital forensics, where it is not possible to predict the method or outcome of recovery. Up until as recently as 2014, very little was known about the destructive functions of SSDs such as garbage collection and the TRIM function, potentially leading to self-corrosion. In the past few years more and more cases involving forensic recovery of SSDs have occurred, allowing some statistical evidence to be collected. However, little is still known about how to prevent permanent data loss from these storage devices. Before addressing this issue, it is important to understand the SSD's functions and the current methods used to recover data from them (Gubanov and Afonin, 2012).

When embarking on the forensic investigation of today's SSDs, the only assumption that can be made is that the investigator can gain access to data stored on the disk. Any data that someone has attempted to delete, by formatting the SSD, is likely lost forever due to the unpredictability of potential recovery process. Modern SSD functions have made data deletion more efficient for the user, but have hindered the forensic recovery process, especially the quick formatting command which will delete data in minutes. There is also no way to stop this command once it has been initiated, even if the computer is turned off, the command will continue as soon as it is switched back on.

These issues are caused by the increased pace in technological development of the SSD's controller and memory throughout multiple companies, meaning that it is increasingly difficult for forensic investigators to know the correct methods to use for any given SSD as they all differ technologically.

3.5.1 SSD controller

Each SSD contains an embedded processor called a controller. Its main function is to be the bridge between the NAND memory mechanisms and the computers using firmware level code. It is an important component that contributes greatly to the SSDs performance levels as it is responsible for the implementation of many tasks including reading and writing cache, encryption, garbage collection, error detection and wear levelling.

Another of the controller's key functions is to map out bad memory cells and locate spare cells for new data to be written to. Therefore, each time the computer needs to read or write

data, it communicates with the controller and the controller translates that to the SSD, so the function can be carried out.

3.5.2 Writeblockers

During a forensic investigation it is crucial that the integrity of the evidence gained is intact. Therefore, digital forensic investigators use devices called writeblockers to gain read-only access to storage devices without compromising their integrity. By blocking any write command made by the host computer to a drive, it protects the data chain of custody by proving the investigator has not tampered with the device.

However, this project aims to explore whether writeblockers have any effect on the implementation of garbage collection and if they can subsequently help prevent data loss. Many papers have dismissed the use of write as an aid to protecting data from deletion during garbage collection. This is because the block deletion programme is run from the drive rather than the host computer, seemingly rendering the writeblocker ineffective.

Gubanov and Afonin (2012) categorically stated that writeblockers have no effect on the functionality and implementation of garbage collection and TRIM functions, with both continuing to perform in the background when a writeblocker is attached. However, it the details of why they do not work are not necessarily explained, leading to further questions about the validity of these claims.

Studies such as Bell and Boddington (2010) tested the theory that the presence of a writeblocker may be able to interrupt the implementation command for the garbage collection and TRIM functions. They suggested that some computers may send the triggering signals via the SATA channel to the SSD, which in the presence of a USB writeblocker may be interrupted and ultimately disabled.

This theory was tested by Bell and Boddington (2010) in a number of experiments. Unfortunately, the outcomes were inconclusive as they were unable to determine the reason for the data alterations which occurred with proof. However, it is stated that there are many different types of SSDs currently on the market, each with varying implementation systems for their garbage collection and TRIM functions, therefore, it is possible that with the write combination of equipment and techniques, writeblockers may be able to assist in the forensic recovery processes.

3.5.3 Garbage collection and TRIM function

While the garbage and TRIM function play a big part in performance and functionality for SSD, they propose tricky advances when data needs to be recovered for forensic investigations. As the garbage collection performs while the computer is idle, it is possible that a forensic investigator could lose potential evidence without knowing because the garbage collection is taking place. Another issue can occur when garbage collection operates during an investigation, the mathematical hash of the image file could change, threatening the validity of the evidence in court.

3.5.3.1 TRIM function

On HDD storage it is possible for data to be overwritten onto the same location as historic data. For SSDs this is not possible, instead new data must be saved to a brand new location on the drive and any pages with old data on them will be deleted entirely before they can be rewritten. The operating system of a computer will detect the logical locations of data and mark it as free. But the computer does not know where the physical location of the file is, this is down to the SSD controller, in particular the TRIM function.

The TRIM function is described as a low-level command that is sent from the host computer to the SSD to inform the which block require deletion. The function allows the computer to function faster and more efficiently as the SSD is able to delete information independently while the computer is idle.

The TRIM function has been supported within Microsoft Windows since Windows 7 and is compatible with any motherboard that has an ATA socket. The ATA is the connection point between the storage device and the main computer which accommodates the commands between the two.

There are three different types of TRIM that are defined by the SATA protocol and appear in different types of SSDs, they are as follows:

- Deterministic Read Zero after TRIM (DZAT)
- Deterministic TRIM (DRAT)
- Non-Deterministic TRIM.

A few years ago, SSDs were only occasionally supported by DRAT, whereas now almost all models come with DRAT or DZAT. DZAT TRIM works by zeroing all read data after the TRIM function has worked until new data is written to the space. The DRAT TRIM will return the same data or become determinate, until garbage collection occurs, or new data is written to the space. Both of these types of TRIM are described as deterministic as there is a known outcome for the read data after TRIM has been applied. However, the Non-Deterministic TRIM works differently as the data returned after this TRIM has been applied may be different to the data before.

For the forensic investigator, DRAT TRIM tends to be the most helpful form of TRIM as it is the most likely to retain the deleted data before new data is written in its place, allowing a period of recovery time. Non-deterministic TRIM has the potential to do the same, however the return is unpredictable and could result in misleading information that could threaten the validity of evidence. DZAT is also unhelpful to a forensic investigator as it returns only zeros, which can be counted as useless information. At present there is limited literature that examines the differences between the three TRIM types. Instead, much of the existing information can be found on websites and other unreliable sources. Therefore, more research is required in the field to define the actual functions and outcomes of each.

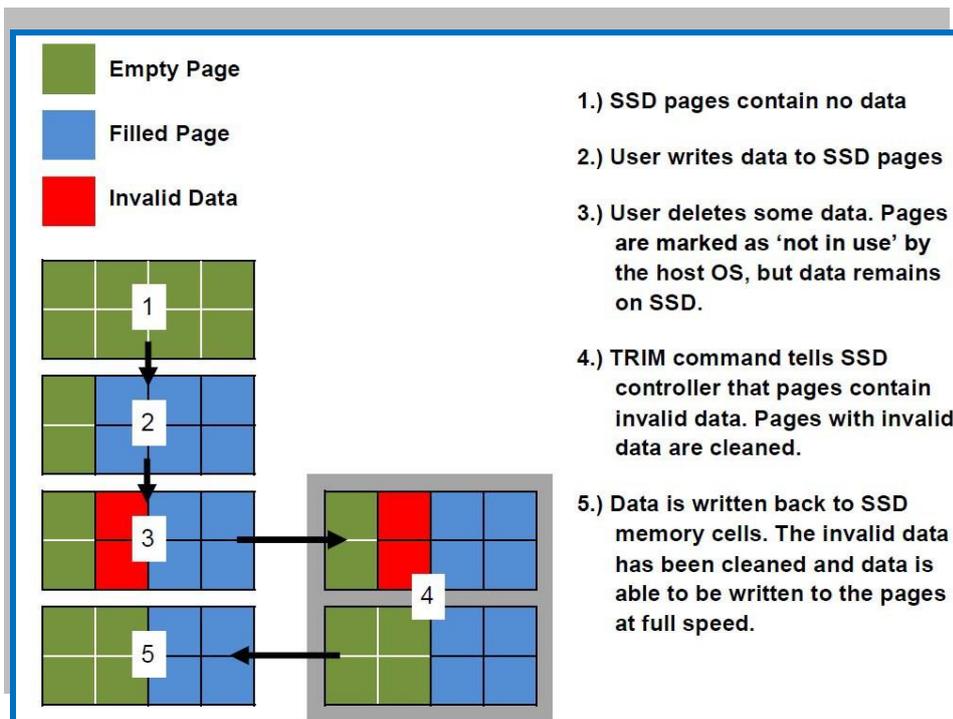


Figure 1: Impression of the TRIM function processes (Gubanov and Afonin, 2012).

3.5.3.2 Garbage collection

The predominant difference between HDDs and SSDs is that SSDs use NAND flash memory to store data, instead of the magnetically coated metal platters within HDDs. This is significant as NAND flash memory does not allow data to be saved and overwritten to the same location. Instead the data must be deleted before new data is saved to that same location. This process of moving data to a new location and deleting the old data is called garbage collection.

To understand this in more detail it is important to understand how flash memory operates. Flash memory is composed of blocks, these are then segmented into smaller sections called pages. Data can be written onto these pages, but it is not possible to delete a single page. Therefore, to delete data from pages to recover space, all valid data must be copied to another location to allow the entire block to be deleted. The clear pages within the deleted block will then be ready to receive new data. Figure 1 demonstrates the process of garbage collection within a simplified image.

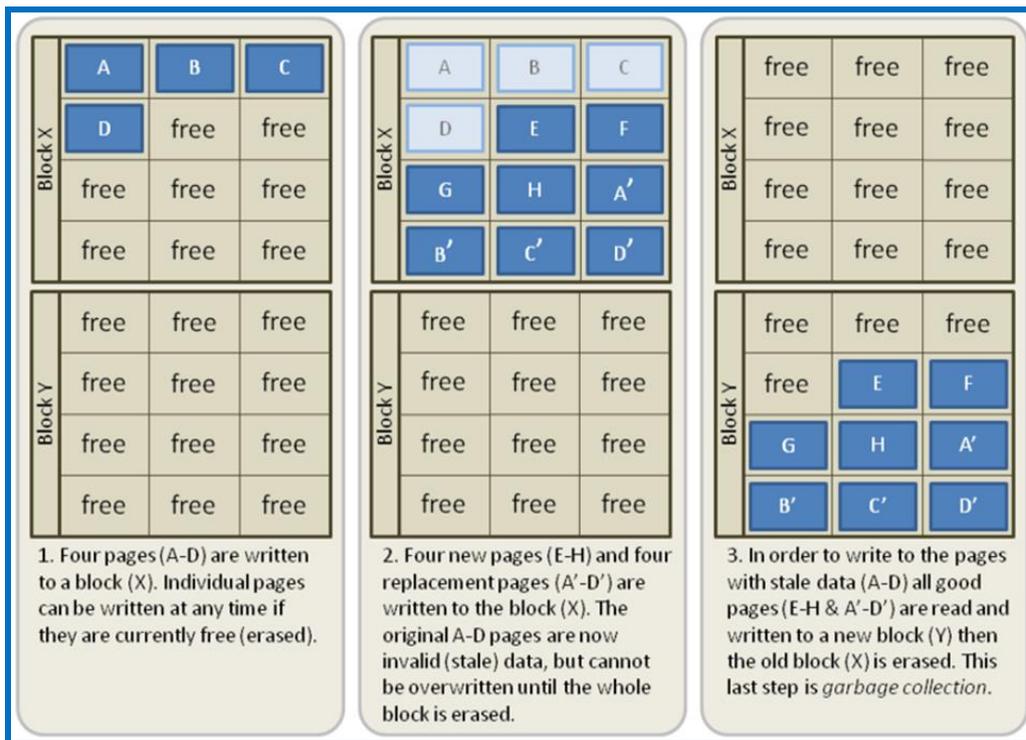


Figure 2: Impression of the garbage collection function (Thatcher, 2009).

When the SSD became popular amongst large numbers of manufacturers, the competition to create the most effective storage device increased. This has resulted in a market which lacks standardisations across the field. Manufacturers have concentrated on the development of the garbage collection function in particular, making the function more efficient. In doing this, they have altered the way in which the function is implemented for different SSDs. The continuing competitive designs of SSDs have lead manufactures to become secretive about their products to ensure they stay competitive.

Although this is good for the markets and developing innovating technology, the secrecy and constant change that's occurring within the design of SSDs is creating challenges for forensic investigation. If forensic investigators cannot work out how the hundreds of differing algorithms work for all of the different SSDs, then it is almost impossible for efficient forensic investigation to occur on these modern devices. This is worsened by the manufactures withholding information as this means that the algorithms that dictate the functionality of garbage collection and TRIM cannot be studied further.

3.5.4 Project Discussions with Gwent Police

On the 21st of February 2018, the first correspondence email from Gwent Police, was received stating the problems faced by their digital forensic investigators. He outlined the issues they were experiencing whilst attempting to investigate and recover data from SSDs, in particular, the M.2 models. He also described what equipment they are currently using, this included:

- A Tableau TDA7-2, which is designed for M.2 SSDs. Connected to the workstation via a Tableau T7u Bridge.
- An M.2 to SATA caddy, designed for use within a laptop that only has SATA connections. Connected to a Tableau T35u.
- An M.2 external caddy that allows it to be used like a portable hard drive. This was connected to a workstation via a Tableau T8u and connected directly to a workstation running CAINE 8.0.
- Booted into a live CD of CAINE 8.0 on the target machine.

The equipment provided by Cardiff University was similar to, but not identical to that used by Gwent Police Force

3.5.5 ACPO Good Practice Guide for Digital Evidence

At present the standard guidelines for digital forensics for the Police Forces of England, Wales and Northern Ireland is the ACPO Good Practice Guide for Digital Evidence. This document provides instructions and regulations on how to gather digital forensic evidence that can be presented in a court of law. It contains information on capturing, analysing and presenting high quality digital evidence (Digital-detective.net, 2018).

Although these guidelines cover a vast range of forensic investigation types, its major flaw is that it has not been updated since 2012. When referring to the timeline of computer storage, it is clear to see that technology has rapidly advanced since 2012 with SSDs becoming the increasingly popular means of storage. In fact, there is no distinction between HDDs and SSDs mentioned within the document. This has caused considerable issues for police forces as there exists no detailed guidance on how to deal with the forensic challenges posed by SSDs.

By overlooking the change in computer storage facilities, there are a number of issues that the police have struggled with, especially relating to the functions of SSDs. The functions of garbage collection and TRIM have already been discussed within this paper, it is known that they are used as within SSDs as an efficient way to create as much storage space as possible, reduce wear and delete unwanted, unnecessary data. However, these functions have also been known inhibit the full forensic recovery of data. There is no advice in the guidance notes to address this issue, instead they simply offer the same set of instructions regardless of the type of hard drive it is and functions it has.

When an SSD is formatted by a perpetrator before the police begin forensic recovery, the action of formatting will continue as long as the computer is switched on. This causes a huge problem for both the recovery effort and the validity of the final evidence gained. Firstly, because as soon as the device is turned back on the formatting will continue and destroy valuable evidence, but also the MD5 hash, the function that ensures the integrity of the data

on the device, will change. This is significant as the MD5 hash is one of the only ways in which a court can tell if digital evidence has been tampered with or not. So, if the police cannot provide the court with legitimate evidence that shows a matching pair of MD5 hash then it may not be considered as valid.

Due to the increase of competitiveness between manufacturers, little data is released regarding the mechanisms behind the garbage collection and TRIM functions within modern SSDs. This makes it even more difficult for the police to produce standard guidelines for forensic recovery of SSDs across the board as each manufacturer's product may require different techniques to recover data from. In some serious cases, police have had to contact manufacturers to request data recovery from one of their SSDs. This is obviously a laborious and unsecure way to gain access to data, which may threaten the integrity of the drive and confidentiality of the victims involved.

The guidance outlines the steps taken to obtain the suspects devices and recover any data that is stored on them. In the section explaining how to capture a device it suggests that one must take care when dealing with volatile data on live systems. This guidance relates mainly to RAM and any online activity that is ongoing when the police arrive on scene. It is suggested that the investigator records any actions they take to obtain any volatile data and track any changes made to the evidence so as much of it remains as valid as possible. This is sound advice when dealing with RAM, however, if this was applied to a situation where an SSD had been formatted then a completely different set of instructions would have to be issued to deal with that situation.

When describing the best practice for analysis of obtained data, there is no details provided on how data would be extracted and analysed from any device, instead it is just stated that a 'forensic strategy' is required to extract only the data relevant to the case. When referring to SSD recovery, it is impossible to assume that the entirety of the data stored on the drive will be available for extraction by the forensic investigator. This is because once it is retrieved from the suspects house, if a quick format has been initiated, then as soon as the device is switched back on, there is no way to stop the destruction of data.

These examples show that the guidelines are in urgent need of updating, along with help from the manufacturers to assist the police in their digital forensic techniques. If official guidelines on the proper ways to recover data from each SSD type became known to the police force, then they would experience far more success in gaining increased amount of data and valid evidence to present in court.

3.5.6 Recent Studies

In recent years there have been a few studies that claim they found solutions for the inconsistency of data acquisition from SSDs. One study completed by Mitchell *et al.* (2016), of Middlesex University, proposed a method called deconstruct and preserve (DaP). This method was not designed to aid recoverability of lost data, but instead focuses on preventing garbage collection from contaminating the device. The paper states that DaP is suitable to use on a variety of devices and results in identical MD5 hashes.

The method is based on the following six key steps that they suggest will become the new Standard Operating Procedure (SOP) for imaging SSDs:

1. **Pre-verification** – This involves providing a hash of the SSD
2. **Deconstruction** – Finding the location where the boot sector is stored, commonly either the MBR or VBR. Once located, the offset of the start and end address is recorded. This is then extracted and moved to a location which is secured by the custodian.
3. **Preservation** – This requires the inserting of MBR into zeros and inserting an exclusive XOR to make it stable.
4. **Acquisition** – Make a duplicate of the original drive.
5. **Reconstruction** – Request offset of the start and end address record from the custodian and insert an exclusive XOR the to preserve the SSD, finally create an image of this.
6. **Verification** – Create a hash of the drive to ensure it matches the original hash.

The purpose of this procedure is to ensure consistency and stability of images taken from SSDs. However, although the theory holds potential for a breakthrough, in practice it may not be appropriate for use in the police forces. Steps 2 and 3 both require a competent DaP practitioner to complete them properly. This would require a full, in depth training for the entire forensics community and the police digital forensics teams, which would prove costly and time consuming.

Additional issues arise if a digital forensics officer was required to take evidence at the scene of the crime. This is not only because the officer requires specialist DaP training, but the steps outlined in the study do not provide enough clarification on how to implement DaP in the variety of real life scenarios that occur. For example, there is nothing mentioned relating to finding a computer in a crime scene that is switched on. This vagueness does not instil confidence in this method, but the study may have provided a starting base from which to continue investigation of this kind.

Another known method of performing data acquisition on an SSD is called chip-off forensics. This is an invasive forensic acquisition procedure that entails physically removing the NAND flash memory chip. It requires an expert to safely remove the chip, with extremely specialised tools. From this the further expertise is required to create images of the chips contents and then fully restore the data through manipulation to provide the coherent contents of data.

This method is effective in retrieving data, post formatting, even if garbage collection and TRIM functions have been active. However, the techniques for retrieving these chips are extremely difficult, requiring desoldering and then cleaning the epoxy off them. This procedure is very delicate and if done incorrectly could irreversibly damage the chip and its contents, making it impossible to acquire any data at all.

The level of expertise required to carry out the removal and recover of data from the chip-off method maybe beyond the capacity of many digital forensic investigators and Gwent Police force. The expense for training individuals to this level and providing them with the particular

equipment they require would take a matter of years. Therefore, it is unrealistic to assume that this method will become SOP throughout the forensic community and police forces.

3.5.7 Conclusion for Background

The research conducted for the background information shows that there has been a long and complex history to computer storage, with rapid developments within the last decade. Although this has benefitted the users of such devices, recent advancements of SSDs have hindered the abilities of digital forensic investigators.

It was also realised that the current issues relating to garbage collection and TRIM functions lack any major studies into their implementation and effects upon data deletion. Therefore, during the research for this project, the project outcomes and scope have slightly altered to ensure the tasks set are not too ambitious for the allotted time scale.

The main points gained from the background research include:

- The issues are ongoing and continually developing
- Manufacturers of SSDs are not willing to openly offer information about their SSDs functions
- There are some existing theories about how these issues can be solved but many of the studies are conflicting and none offer a guaranteed solution.
- TRIM triggers the embedded processor in and SSD to deploy garbage collection
- And the most up to date guidelines used by the UK police (ACCPO) are redundant when working with modern SSDs.

4. Specification & Design

4.1 Approach

4.1.1 Details of Experimental Equipment

For this project the experiment to test the effects of garbage collection during the forensic process of data extraction from an SSD, required the following equipment, all of which was provided by Cardiff University:

- M.2 SSDs (x3)
- Forensic bridge/Writeblocker
- Forensic Tool Kit – Forensic recovery software
- Image of a hard drive
- USB to SATA and IDE Adapter
- M.2 SSD to 2.5in SATA SSD Converter

4.1.2 Manufacturing Information and Details of Each SSD

To ensure the experiments are unbiased and to increase validity three SSDs, sourced from different manufacturers, were tested instead of only one type. Another reason why it is important to test a variety of models is due to the different ways in which manufacturers implement the garbage collection mechanism. This means that tests may be able to reveal the subtle differences between the effects of garbage collection on data recovery across different manufacturer's products, providing valuable insight into which models forensic investigators can recover more data from. Other benefits of using multiple SSDs is to test if the data recovery from any models are affected by the presence of a Writeblocker.

Prior to the start of any experimentation, all manufacturers were contacted for research purposes in an attempt to gain any additional information on the details of garbage collection implementation for their particular model of SSD. Understandably, each company stated that they preferred not to divulge details on the individual SSDs if there were to be revealed in a research paper, as these were crucial to the competitiveness of the company's products. Therefore, the following information on each SSD is sourced from public knowledge of the product and any additional information the companies were willing to have published.

1. WD green pc Sata M. 2 2280 SSD

Attributes:

Cost – £54

Capacity – 120 GB

Height – 1.5mm

Read Speed – 540MB/s

Write Speed – 430MB/s

This SSD model was the cheapest out of the three tested, but it had extremely limited information available both in the manual and online about its garbage collection function. In an attempt to gain further information, the company was contacted but they did not respond.

2. Crucial MX300 M.2 SSD

Attributes:

Cost – £77.99

Capacity – 275GB

Read Speed – 530MB/s

Write Speed – 510MB/s

Online is states that this SSD used active garbage collection and TRIM support functions, however, not further details were provided. When the manufacturer was contacted, replied with the following statement:

“Hello Michael,

Further to our earlier chat I have now had a reply from our engineering team.

Whilst they would be happy to call you to discuss this I need to advise you that they would only provide industry standards when it comes to the information that you are looking for.

The specifics to what you are requesting are something that we can only share if you are happy to sign a non-disclosure form. However, this would be pointless as you need to write a paper on this.

If you wish to discuss this, please provide me your number and I will pass this to my engineering team so that you can discuss this further.”

The above email evidences that the company was willing to discuss the garbage collection feature of this particular SSD. However, this would only occur once a non-disclosure form was signed and none of the information could have been within the project. Therefore, it would be difficult to validate any specific theories based upon the garbage collection features of this SSD as they could not be fully discussed within this paper. This does however reveal that each company implement the garbage collection function in different ways as the privacy surrounding the subject suggests that they wish to retain information that could be taken advantage of by other competitors. “

3. Kingston SM2280S3G2 M.2 SSD

Attributes:

Cost – £52.82

Capacity – 120GB

Read Speed – 550MB/s

Write Speed – 520MB/s

This SSD was priced mid-range and appeared to provide the most information regarding the garbage collection and TRIM function of the product. The website states that this SSD uses an LSI SandForce controller that allows the SSD to implement its garbage collection function more efficiently by preparing blocks for deletion in real time as opposed to others which apply is as a background action. They suggest that real time block preparation maximises the benefits of low power mode, which is induced by powering down the system when it does not need to be accessed. Kingston was contacted and asked if they would provide any information on their garbage collection for my project, the responses are displayed below.

“Dear Micheal,

Thank you for contacting Kingston Technology.

Kingston receives a high volume of requests to assist with projects and studies, although we would like to assist everyone it is unfortunately not feasible.

All information that we make available to the public is on our website on the following link:
<https://www.kingston.com/en/company>.

Kind regards.

Silvia Giglio
Customer Service and RMA Manager”

4.1.3 Forensic Bridge/Writeblocker

For this experiment, the forensic bridge has been used to stop any additional data from being written onto the drive after the forensic investigation has commenced. It has also been used to gauge whether or not it has any effect on the garbage collection initiation.

The forensic bridge used in this experiment was a Tableau Forensic SATA/IDE Bridge T35u. The product is described as an industry standard, portable writeblocker that allows forensic acquisition of IDE and SATA SSDs.

4.1.4 Forensic Toolkit

The forensics toolkit is a computer forensics software created by Access Data used to recover deleted data including pictures, videos and emails. This product will one of the main forensic software used in the digital forensics community.

5. Implementation

5.1 Constraints

Finding solutions to the many issues of SSD forensics is a substantial challenge considering the rapid development of SSD manufacturing. Therefore, there are limitations to the solutions that can be discovered within one project.

One of the main constraints within this project is the lack of a team. Many studies that have been carried out on the topics relevant to this project are conducted by a team of people. With more people bringing their own speciality knowledge to a project, there can be a wider platform for idea discussion and creativity.

When conducting digital forensic investigation, there is a requirement to have all of the relevant equipment, otherwise the outcomes will not reflect the real-life significances. Although Cardiff University provided the essential equipment, it did not necessarily reflect the conditions that Gwent Police use for their forensic investigations.

The SSDs used in the experiments range in size from 120GB to 275GB. These SSDs required imaging and storage multiple times, the laptop used was 256GB, the only viable option was to use an external hard drive that could accommodate this. The external hard drive used was 1TB, which was big enough to store only one experiment but was not large enough for the rest. This meant that some SSD images had to be deleted, after the experiment was finalised. This does not reflect how the scenario would occur during a real life investigation, as there would be a sufficient amount of equipment to store all of the data.

In addition to limited resources and time, a dissertation project also has limited funding. Many companies fund research into similar projects which will inevitably allow the researcher to produce results with the latest technology and equipment. However, it is hoped that this project will provide solutions for some of the unanswered questions asked by Gwent Police Force, to help them during future digital forensic investigation of SSDs as well as form a basis for further research into these areas.

5.2 Experiment Methods

The image used in this experiment was created using a fresh, bootable Windows 10 on an external hard drive. It contained a folder of pictures which was created to imitate the type of files that are commonly recovered during a forensic investigation. The pictures used within the experiments were obtained from a copyright-free website that provides stock photos (Pixabay.com, 2018).

The external hard drive was then copied using the FTK Imager in Raw Image Format (also known as DD). This is one of the main formatting types used in digital forensics because it copies data using a bit by bit stream.

When the FTK Imager creates an image it also creates a MD5 hash. These will be checked each time an image is copied to the SSD to verify the image has not been altered during the copying process.

This main image file was then copied onto the host computer using FTK imager(DD) and the MD5 was checked to see if it matched.

5.2.1 Creating the Image

Each experiment in this project began with an image of the 80GB external hard drive that was under investigation. This image was then transferred to each of the three SSDs, ready for experimentation. The three SSDs vary in size, so it was important to create an image which was able to fit onto every SSD. To ensure the experiments imitated real life scenarios, experienced by Gwent Police, as much as possible, the drive image was made to be bootable.

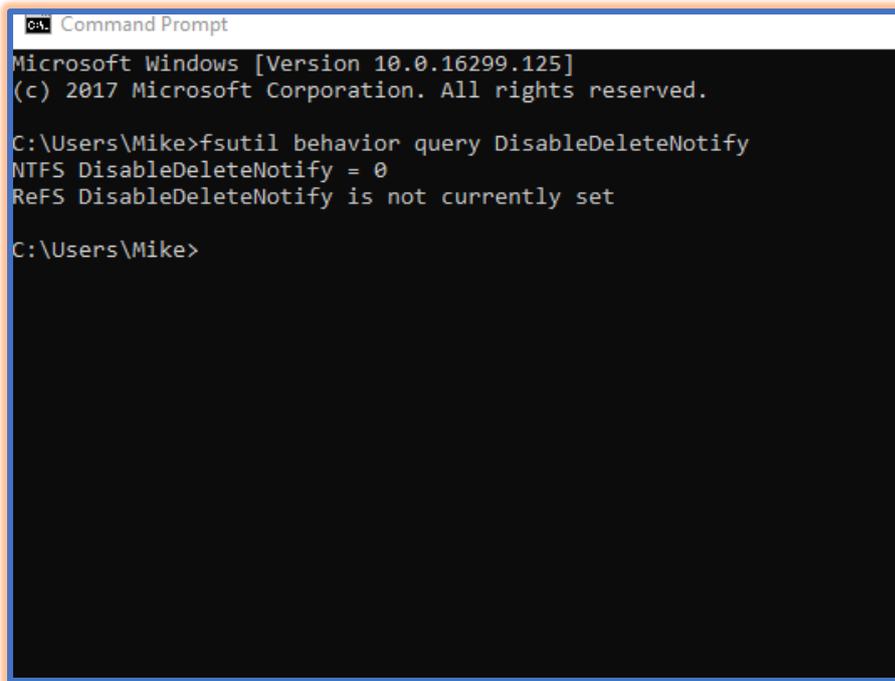
The first steps taken to create the image was to download a fresh version of Windows 10 onto the external hard drive. Initially, it was to be sourced from the university but due to complications the download was unavailable. It was finally sourced from the Windows website, which offered free downloads of Windows 10 for a recovery disk. It is not possible to directly download Windows 10 onto the external hard drive, the program WinToUSB was used as it is specifically designed to mount Windows onto external drives.

To create an image that represented a realistic hard drive under investigation, files were placed on a USB, to act as evidence. A folder containing pictures of cats, with a range of file sizes, was uploaded, totalling a file size of 1.58GB. The WinToUSB was then used to mount the hard drive to boot as a normal Windows computer. The pictures were then placed into the documents folder ready for the experiment. These pictures would act as evidence samples during the forensic recovery experiments by deleting a number of them and then attempting to recovery.

Finally, the hard drive was imaged. The external hard drive was copied using the FTK Imager in Raw Image Format (also known as DD). This is a predominant formatting type used in digital forensics because it copies data using a bit by bit stream. This image was the main one transferred onto all of the SSDs for experimentation using a program called Win32Imager, which provided the same Raw image format.

5.2.2 Experiment Preparation

Before beginning experimentation, equipment and techniques need to be reviewed to ensure the tests are fair and non-bias, to ensure validity of results. The main focus of this project is the TRIM function and garbage collection, therefore all SSDs were checked to see if they supported both of these functions. This was performed by using the command line on with administration privileges and typing the command “fsutil behavior query DisableDeleteNotify”. If the results show the number 1, then this means that TRIM is currently disabled, however, if a 0 appears, then the TRIM will be enabled. To check that each of the SSDs had the TRIM function enabled, they were booted up separately and running this command. A screen show of this is shown in Figure 3



```
Command Prompt
Microsoft Windows [Version 10.0.16299.125]
(c) 2017 Microsoft Corporation. All rights reserved.

C:\Users\Mike>fsutil behavior query DisableDeleteNotify
NTFS DisableDeleteNotify = 0
ReFS DisableDeleteNotify is not currently set

C:\Users\Mike>
```

Figure 3: Screenshot showing that the TRIM is enabled

It is also important to ensure that the data that will be tested on is of similar format to the evidence Gwent Police come across. It was stated that Gwent Police usually deal with cases involving pictures, therefore, it was decided that a file of pictures would be used to test on. 311 pictures were downloaded and then stored on each SSD under the file name "cats". the 311 pictures will be deleted from the folder then recycled from the recycle bin.

6. Results and Evaluation

6.1 Experiment Results

6.1.1 Experiment 1

The first experiment will be used to gauge how effective the presence of a writeblocker is against TRIM and garbage collection functions by checking how many files are recoverable after they have been deleted. This will be achieved by deleting the 311 pictures, which will in turn trigger the TRIM function on the SSD. The SSD will then be immediately shut down.

The writeblocker will then be connected and the drive imaged. FTK imager will be used to check if any of the files can be recovered. If the writeblocker has any effect, the MD5 hash taken at the beginning will have altered from the final hash. The FTK imager provides a feature which tests the MD5 hash beginning and end and compares them. Also if the MD5 hash changed during the imaging process, another image will be taken and then compared to see the difference in recovery.

Experiment 1 Step by Step Methodology

1. Copy main image to SSD using DD
2. Boot up SSD
3. Delete the 311 pictures
4. Shutdown SSD immediately
5. Connect writeblocker and turn on SSD
6. Image SSD and take note of MD5
7. Use FTK to see what files can be recovered
8. Record results

Expected results

Recent studies have stated that writeblocker have no effect on garbage collection or TRIM, even when connected to a forensic bridge. Therefore, it was predicted that no MD5 hash would match and approximately half of the pictures would prove unrecoverable.

Kingston SSD (120GB)

The first SSD to be tested was the Kingston. This is a mid-range SSD compared to the other two used in these experiments. The 311 pictures were deleted from the drive and was then immediately shut down. The image was connected to the writeblocker and the drive was imaged using FTK imager. An MD5 hash was taken at the start of the image process and at the end to verify it had not been tampered with.

```
Experiment1.001 - Notepad
File Edit Format View Help
-----
Information for I:\kingston image\Experiment1:

Physical Evidentiary Item (Source) Information:
[Device Info]
Source Type: Physical
[Drive Geometry]
Cylinders: 14,593
Tracks per Cylinder: 255
Sectors per Track: 63
Bytes per Sector: 512
Sector Count: 234,441,648
[Physical Drive Information]
Drive Model: ASMT 2115 USB Device
Drive Serial Number: 00000000000000000000
Drive Interface Type: USB
Removable drive: False
Source data size: 114473 MB
Sector count: 234441648
[Computed Hashes]
MD5 checksum: 11bed3a6a221821b9cada399c6b03aa8
SHA1 checksum: f66ef78a6c7b35e45dceb2c3d54afa04dc9ea6fc

Image Information:
Acquisition started: Wed May 02 18:31:38 2018
Acquisition finished: Wed May 02 23:14:26 2018
Segment list:
I:\kingston image\Experiment1.001

Image Verification Results:
Verification started: Wed May 02 23:14:26 2018
Verification finished: Wed May 02 23:32:34 2018
MD5 checksum: 11bed3a6a221821b9cada399c6b03aa8 : verified
SHA1 checksum: f66ef78a6c7b35e45dceb2c3d54afa04dc9ea6fc : verif:
```

Figure 4: Kingston SSD Matching MD5 hash

Figure 4 shows that the Process took 4 hours and 32 minutes and the MD5 was the same at the start of the image process and at the end.

When investigating the image, it was discovered that 291/311 images were recoverable. The 20 images that were unrecoverable had all their data replaced with zeros.

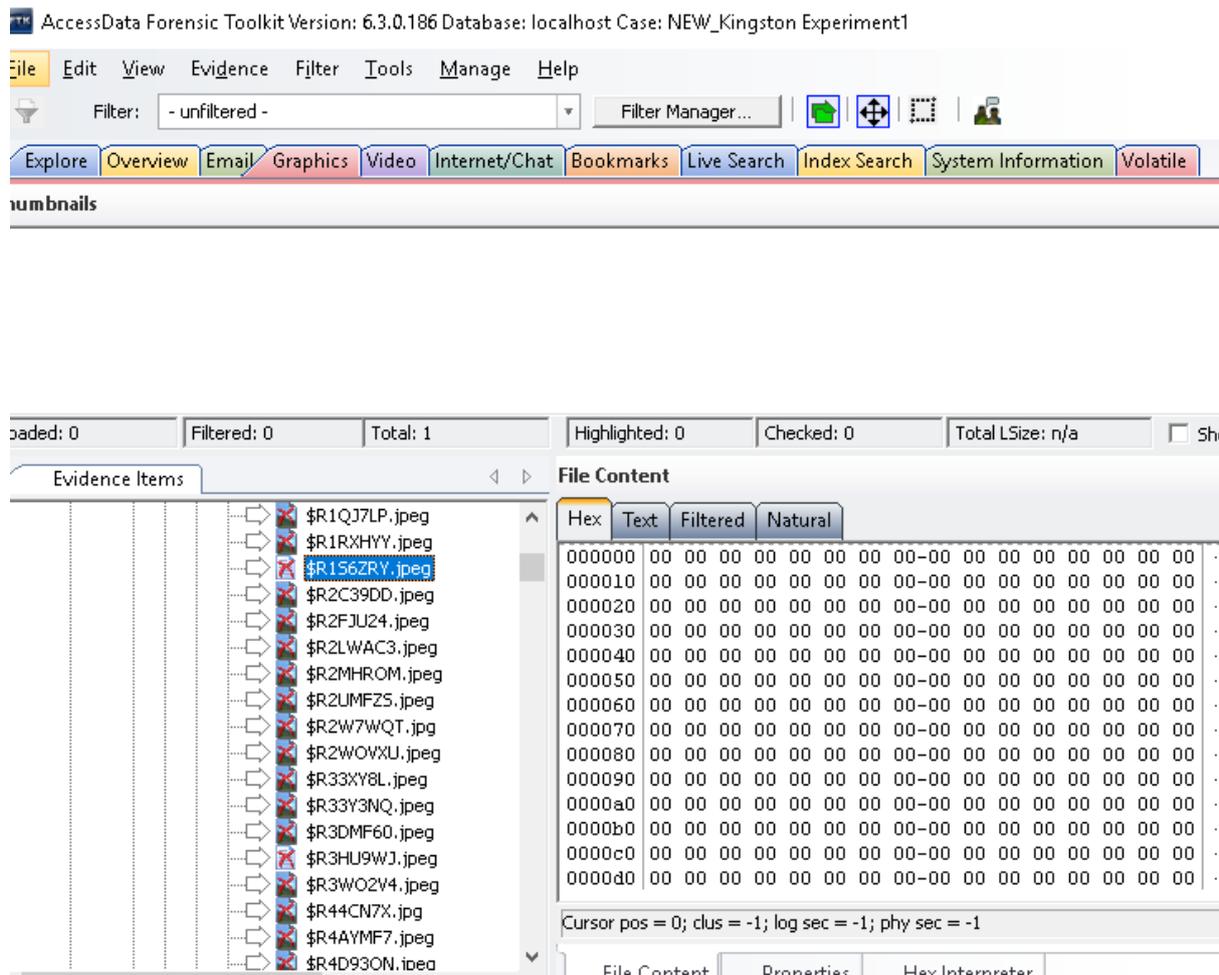


Figure 5: Kingston SSD- files with zeros as data

Upon closer inspection, it was noticed that one of the pictures displayed as much smaller than the original size. It was unclear why this occurred, as all of the data was still intact.

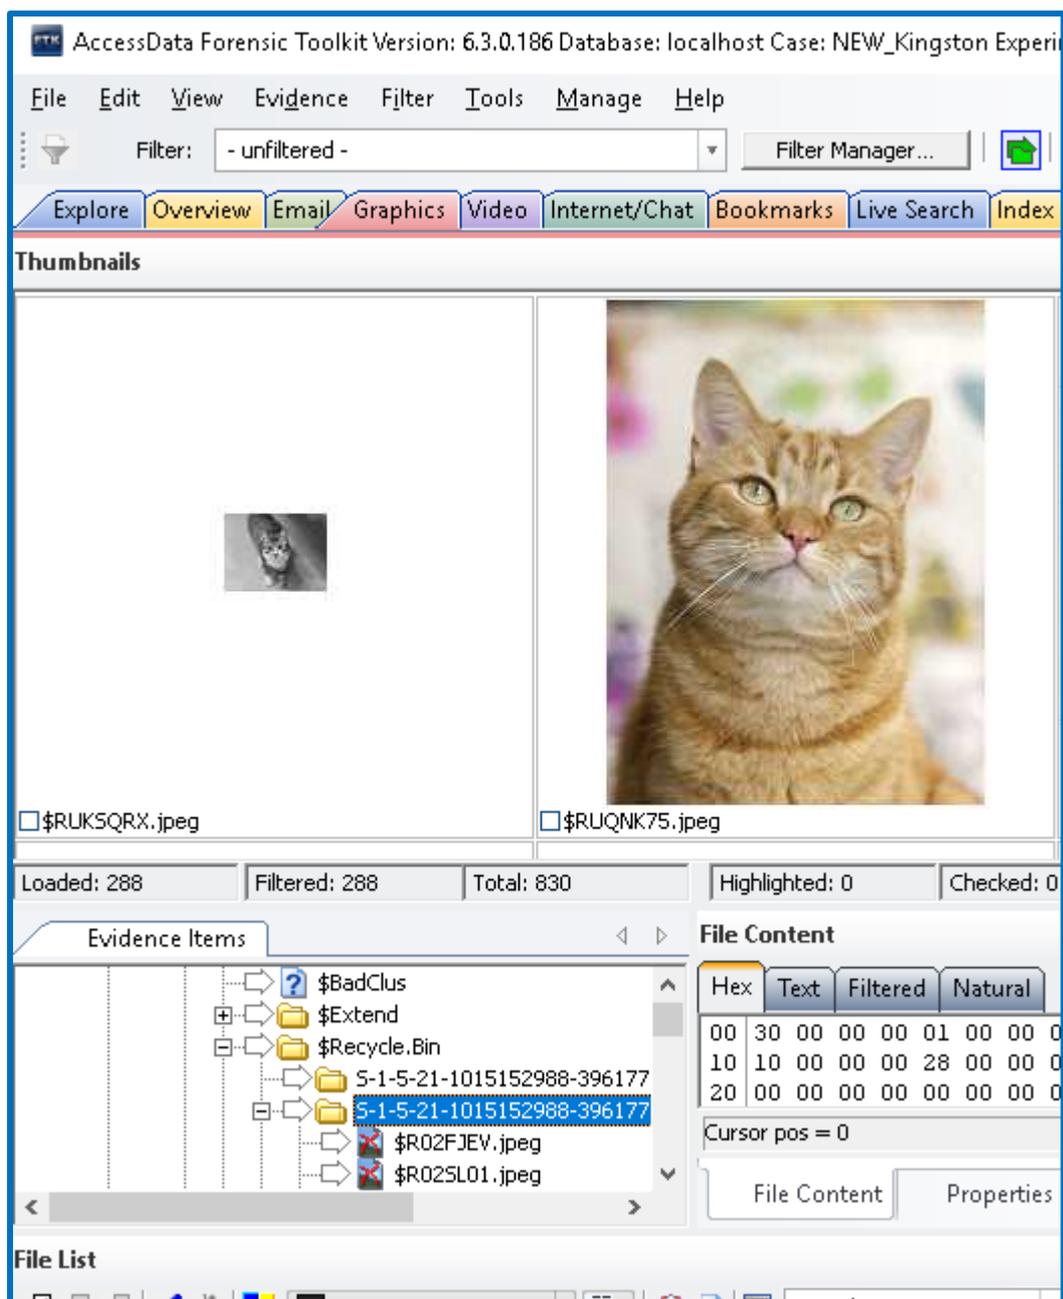


Figure 6: Kingston SSD- Shows picture changes when recovered.

WD Green (120GB)

The second SSD for this experiment was the WD Green. This was the cheapest SSD in experiment and underwent the same process as the previous SSD in the experiment.

```
Experiment1.001 - Notepad
File Edit Format View Help
Information for I:\WD Green image\Experiment1:

Physical Evidentiary Item (Source) Information:
[Device Info]
Source Type: Physical
[Drive Geometry]
Cylinders: 14,593
Tracks per Cylinder: 255
Sectors per Track: 63
Bytes per Sector: 512
Sector Count: 234,441,648
[Physical Drive Information]
Drive Model: ASMT 2115 USB Device
Drive Serial Number: 00000000000000000000
Drive Interface Type: USB
Removable drive: False
Source data size: 114473 MB
Sector count: 234441648
[Computed Hashes]
MD5 checksum: e2a45d96d9f6f3cedc10bed2c5982550
SHA1 checksum: 1ffd7d9c97005b9f0df083a5d62fa9d9f8519f4f

Image Information:
Acquisition started: Thu May 03 00:07:05 2018
Acquisition finished: Thu May 03 04:48:53 2018
Segment list:
I:\WD Green image\Experiment1.001

Image Verification Results:
Verification started: Thu May 03 04:48:53 2018
Verification finished: Thu May 03 05:11:00 2018
MD5 checksum: e2a45d96d9f6f3cedc10bed2c5982550 : verified
SHA1 checksum: 1ffd7d9c97005b9f0df083a5d62fa9d9f8519f4f : verified
```

Figure 7:WD Green Time taken and matching MD5

Figure 7 shows that the Process took 3 hours and 57 minutes and the MD5 was the same at the start of the image process and at the end.

Like the previous SSD there were files recoverable and other files unrecoverable and filled with zeros. The number of recoverable files was 287/311.

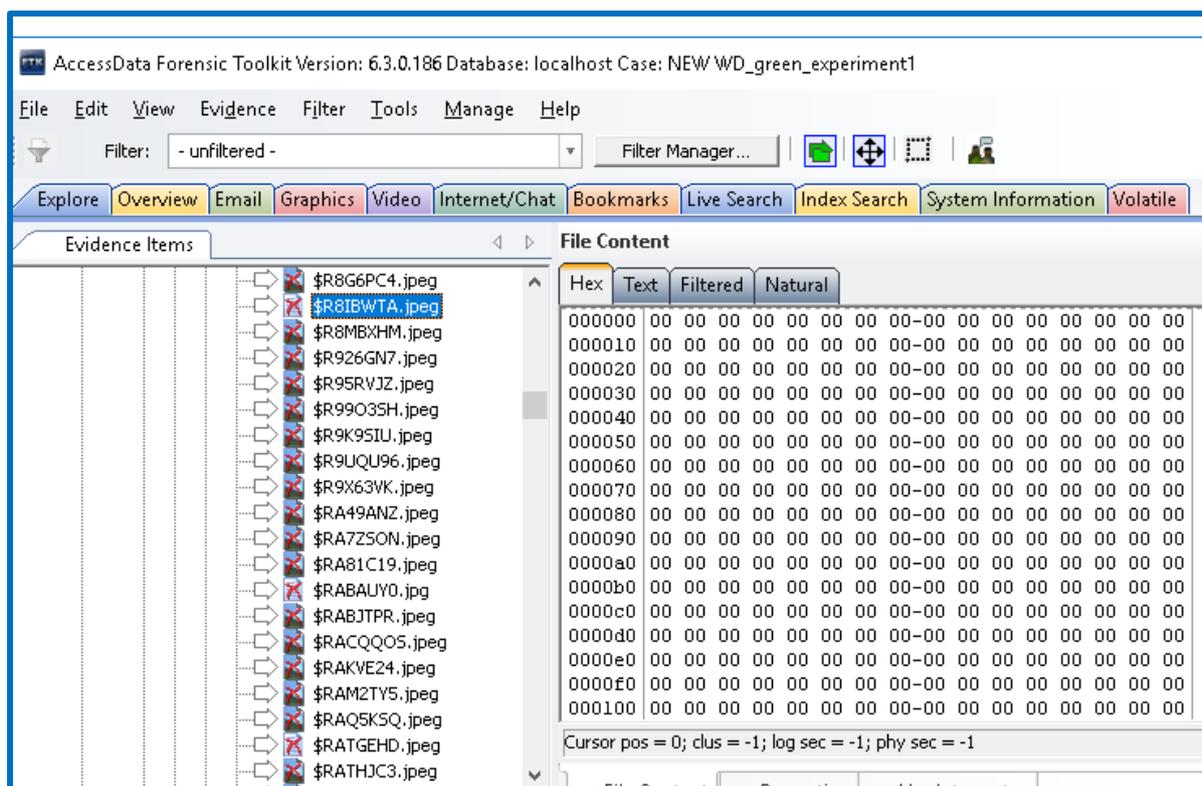


Figure 8: WD Green Files been filled with zeros

Crucial SSD (275GB)

The last SSD for this experiment was the Crucial SSD. This was the most expensive and the largest SSD in the experiment. This SSD went through the same process as the previous SSDs in the experiment.

```
experiment1.001 - Notepad
File Edit Format View Help
-----
Information for I:\crucial image\experiment1:

Physical Evidentiary Item (Source) Information:
[Device Info]
Source Type: Physical
[Drive Geometry]
Cylinders: 33,441
Tracks per Cylinder: 255
Sectors per Track: 63
Bytes per Sector: 512
Sector Count: 537,234,768
[Physical Drive Information]
Drive Model: ASMT 2115 USB Device
Drive Serial Number: 00000000000000000000
Drive Interface Type: USB
Removable drive: False
Source data size: 262321 MB
Sector count: 537234768
[Computed Hashes]
MD5 checksum: 903896aa0c5dcb0b383b9d7fe2661ad7
SHA1 checksum: b89fc60f364b01f27a15da7f6e28c1750c06a81e

Image Information:
Acquisition started: Thu May 03 05:05:21 2018
Acquisition finished: Thu May 03 15:54:09 2018
Segment list:
I:\crucial image\experiment1.001

Image Verification Results:
Verification started: Thu May 03 15:54:09 2018
Verification finished: Thu May 03 16:40:00 2018
MD5 checksum: 903896aa0c5dcb0b383b9d7fe2661ad7 : verified
SHA1 checksum: b89fc60f364b01f27a15da7f6e28c1750c06a81e : verified
```

Figure 9: Crucial SSD Time taken and MD5

Figure 9 shows that the Process took 3 hours and 57 minutes and the MD5 was the same at the start of the image process and at the end.

Like the previous SSDs there were files recoverable and other files unrecoverable and filled with zeros. The number of recoverable files was 294/311.

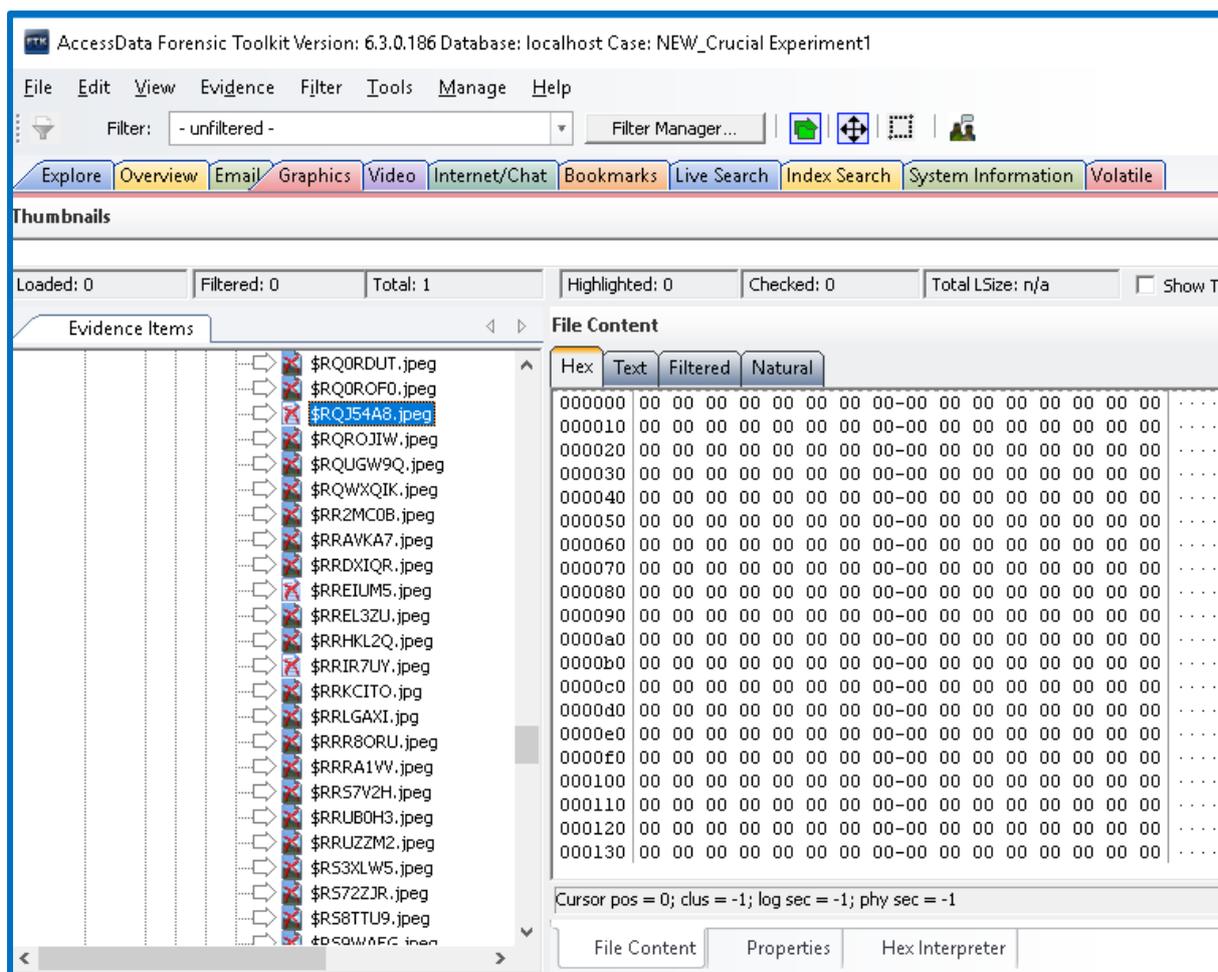


Figure 10:Crucial SSD files been filled with zeros

Actual results

SDD	Did the MD5 Match	How Many pictures were recovered from first image	Time taken to complete image	How many pictures were recovered from second image
Kingston SSD	Yes	291/311 pictures recovered	4 hours and 32 mins	N/A
WD Green	Yes	287/311 pictures recovered	3 hours and 57 mins	N/A
Crucial SSD	Yes	294/311 pictures recovered	10 hours and 48 minutes	N/A

Table 1: Results from experiment 1

Summary of Experiment 1

After conducting the first experiment, the results showed that the Crucial SSD performed the best recovering 294/311 images and the WD green performed marginally worst with 287/311. This could be due to many factors that this experiment did not cover. The Crucial SSD is the largest of the three, therefore, it would have taken more time for files to be zeroed out. Another interesting factor was that if an image took a relatively long amount of time, then more pictures were recoverable.

Some interesting discoveries were made during experiment 1 that can apply to the following experiments. These include:

- The garbage collection for each SSD ran at different speeds
- The MD5 hash stayed the same, regardless of being connected with a writeblocker for 4+ hours. This may suggest that the garbage collection doesn't work while connected to a writeblocker.

6.1.2 Experiment 2

The second experiment aims to justify whether the presence of a write blocker has any effect on the SSD garbage collection and trim functionality or implementation. Previous research on this subject suggested that write blockers have no effect on the forensic processing of SSDs. This theory works on the idea that SSDs have their own embedded computer within their chips which works independently of the write blocker functions, therefore bypassing the process altogether. However, there is only minimal research on this subject, so this experiment aims to validate or challenge existing theories to clarify concept. This will provide direction for future study into this area.

This experiment is based on some of the factors that were discovered during experiment 1 and the pre-experiment tests. One of these factors taken into account was the result of experiment 1, where pictures were recovered immediately after deletion and the shutdown of the computer. Therefore, the information learned has been considered whilst planning experiment 2.

During this experiment, all pictures will be deleted and the SSD shut down. However, this time, instead of imaging the SSDs straight away, the writeblocker will be connected and the SSD will be booted and left idle for an hour. After that, the SSD will be shut down and imaged. Then there will be an attempt to recover the pictures. By comparing the results of experiment 1 to experiment 2 it may be determined whether the presence of a writeblocker had any effect on recoverability.

Experiment 2 Step by Step Methodology

These steps will be performed on all three SSDs:

1. Quick format SSD
2. Copy image to SSD using win32imager in DD
3. The SSD will be booted
4. Delete the 311 pictures from recycle bin
5. The computer will be shut down
6. The write blocker re connected
7. The SSD will be booted back up
8. Leave SSD idle for 1 hour
9. SSD will be shutdown
10. Image drive and Check md5 hash
11. Use FTK to see what files can be recovered
12. Record how many pictures were recovered

Unforeseen Issues and Alterations in Experiment 2

Half way through the implementation of experiment 2 it was noticed that the writeblocker had stopped working, which meant that no results were found to be valid. Cardiff University quickly replaced the writeblocker with one of the same brand but not the same model. When the experiment was restarted, with the new writeblocker, problems occurred with the performance of the new equipment. For an unknown reason, the drive would not boot whilst the writeblocker was connected. Evidence of this is shown in Figure 11.

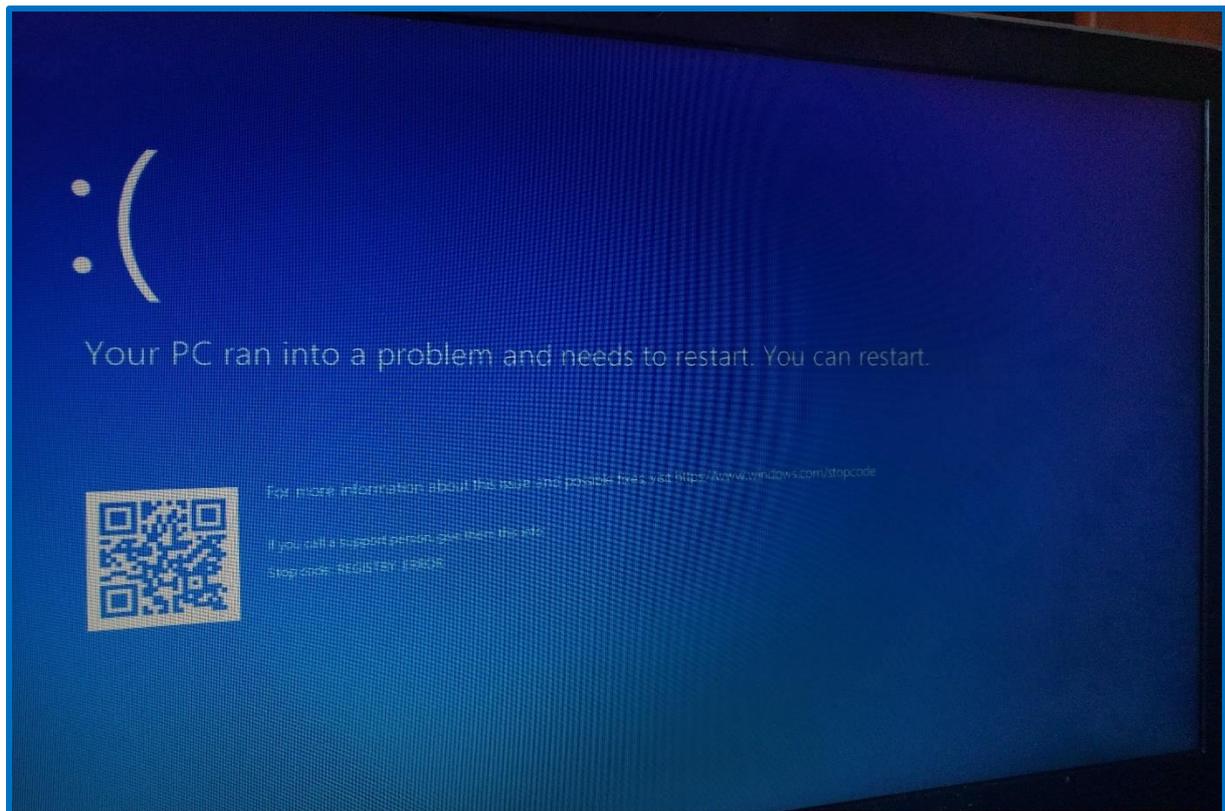


Figure 11: Problems with experiment 2

This caused significant issues for the entire experiment, as it was initially planned to test the difference in recovery when the writeblocker was attached at varying intervals. Therefore, during further research and equipment testing, it was found that whilst the image was being processed within FTK, files appeared viewable. By the time FTK finished processing the files some were unrecoverable, some zeroed and some were text. Further evidence can be seen in the results of experiment 2

New Experiment 2 Step by Step Methodology

These steps will be performed on all three SSDs:

1. Quick format SSD
2. Copy image to SSD using win32imager in DD
3. The SSD will be booted
4. Delete the 311 pictures from recycle bin
5. The computer will be shut down
6. The drive will be imaged and Check md5 hash
7. FTK will be used to see what files can be recovered
8. Record amount of recoverable pictures
9. Record amount of files zeroed
10. Record amount of files that display hex data but no image
11. Record results of files disappearing

Expected results

SDD	Will results be zeroed out	How Many pictures can defiantly be recovered
Kingston SSD	Yes	150/311
Crucial SSD	Yes	150/311
WD Green	Yes	150/311

Table 2: Expected results for experiment 2

WD Green

The first SSD in this test was the WD Green. The SSD will be booted up and the 311 pictures will be deleted. No writeblocker will be connected and the drive will be imaged. Changes within the file will be searched for.

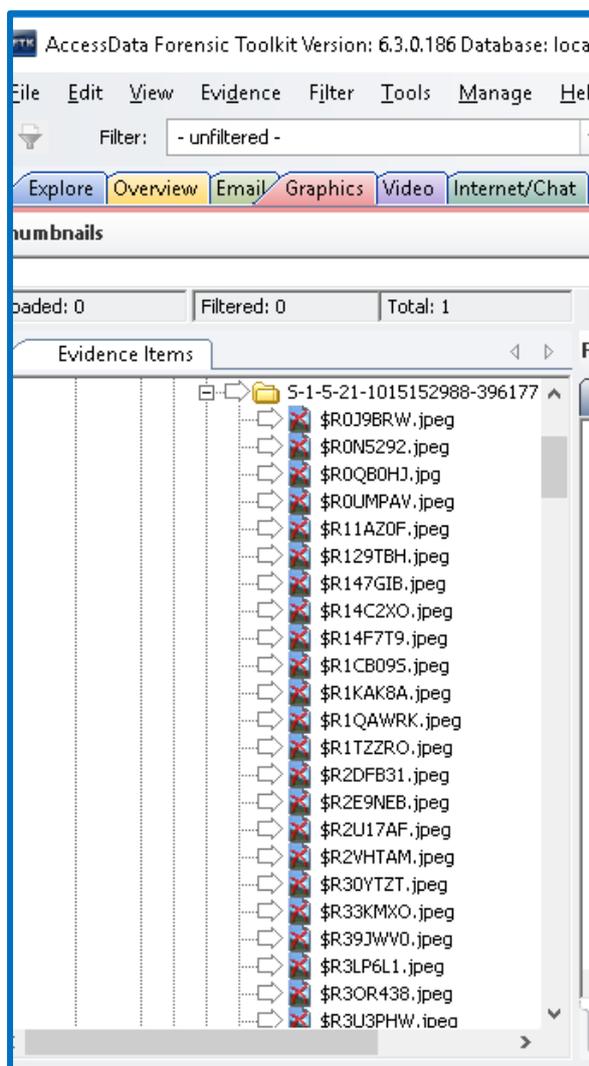


Figure 12: WD Green- Files when first imported to FTK

This is an example of when the image first gets loaded into FTK imager. The majority of the files display images, but the top file displayed text.

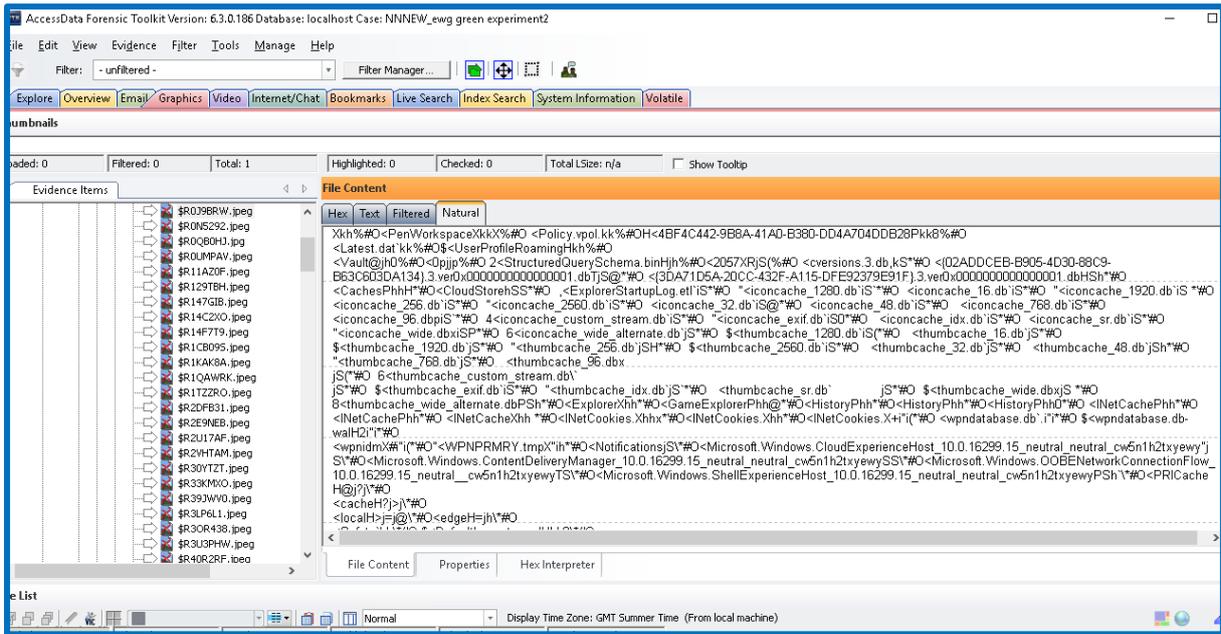


Figure 13: WD Green-File displayed in natural form

It was decided this should be observed for changes, which could provide information about why files were changing while in FTK.

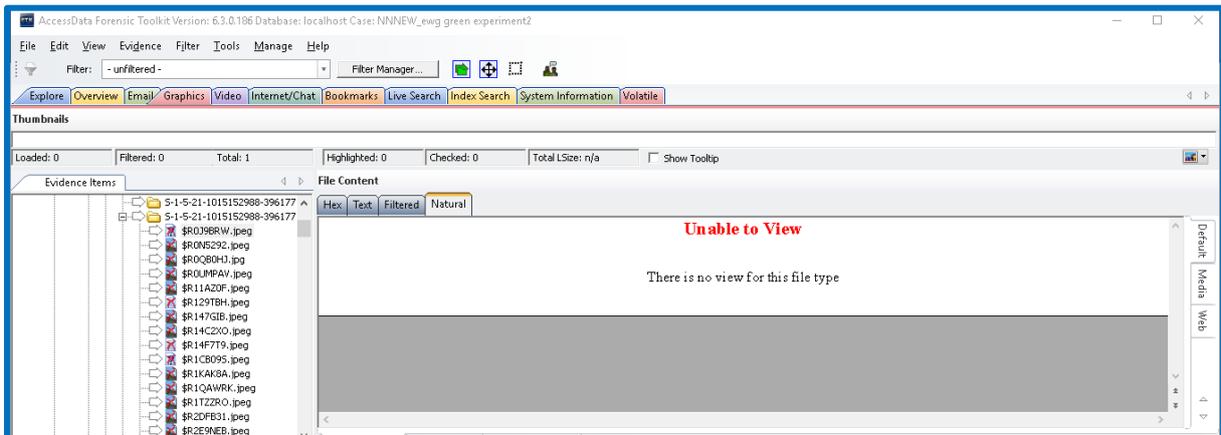


Figure 14: WD Green SSD- Files after being processed

This is a picture of the file when FTK had finished processing it. on closer inspection it was noticed that many of the files that seemed to be recovered at the start had now become unrecoverable. 235/311 were recoverable by the time FTK had finished processing.

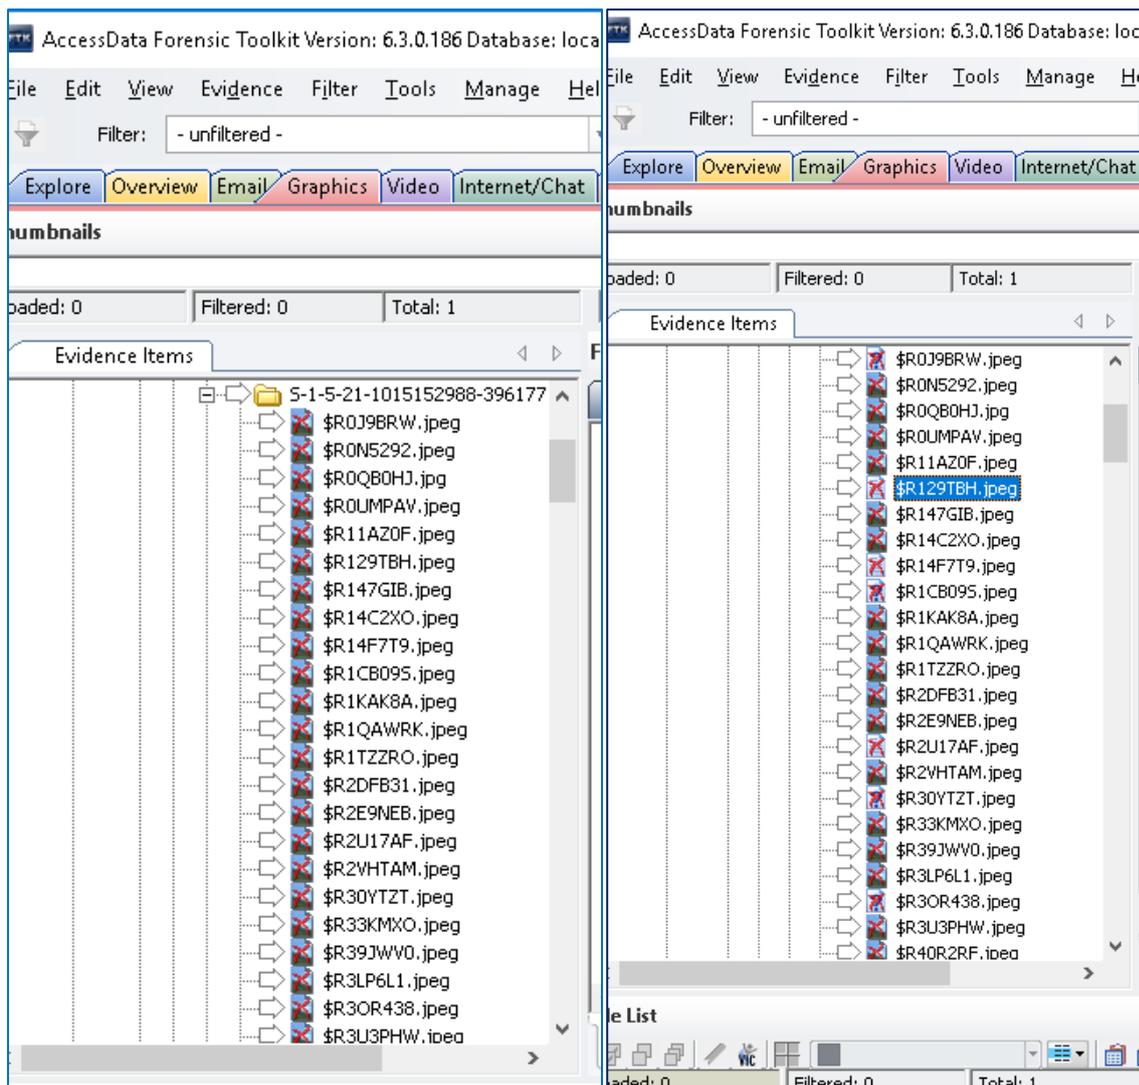


Figure 15: WD Green SSD- Side by side comparison of processed files

This is a side by side comparison of the file list at the start of FTK processing and the end. It shows that 7/23 pictures in the sample were now unrecoverable. At the end of the WD Green testing, it was unclear whether there was data within any of the files that appeared as zeros before they read as unrecoverable.

Kingston SSD

The second SSD to be tested was the Kingston. The same techniques as the first test will be used but with more of a focus on figuring out if there was data on the files before they become unrecoverable. To do this, the FTK was paused when the files were first processed to access whether the pictures were changing on the image or if the data just was not there initially.

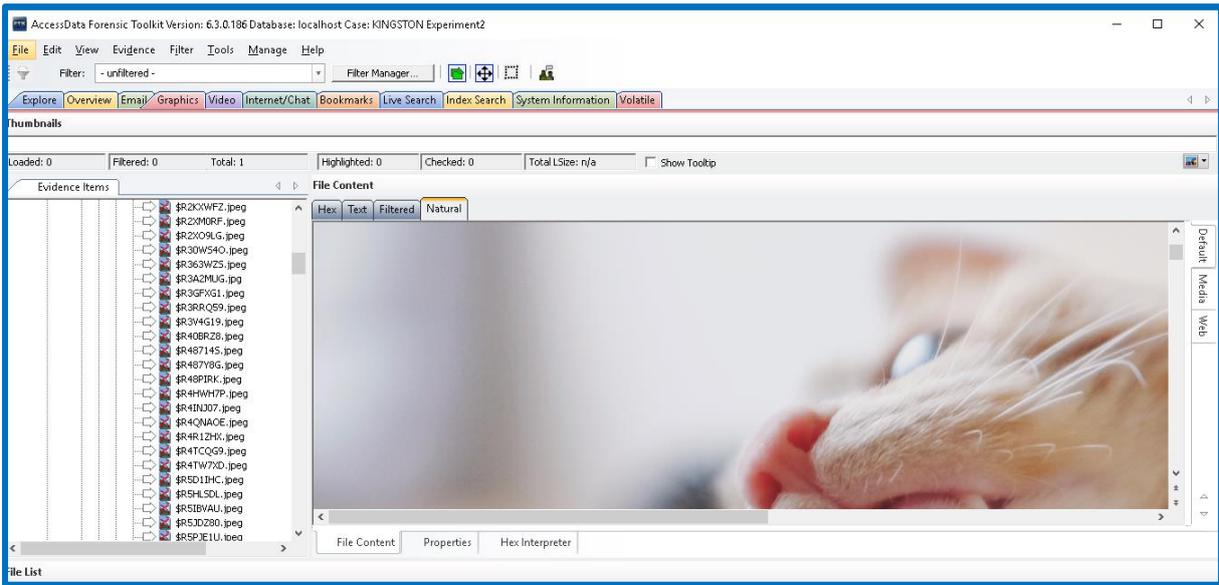


Figure 16: Kingston SSD-After FTK was paused

This is a picture of the FTK paused showing that files were still accessible when paused. The first 25 pictures were selected for examination to test if they were unusual in any way.

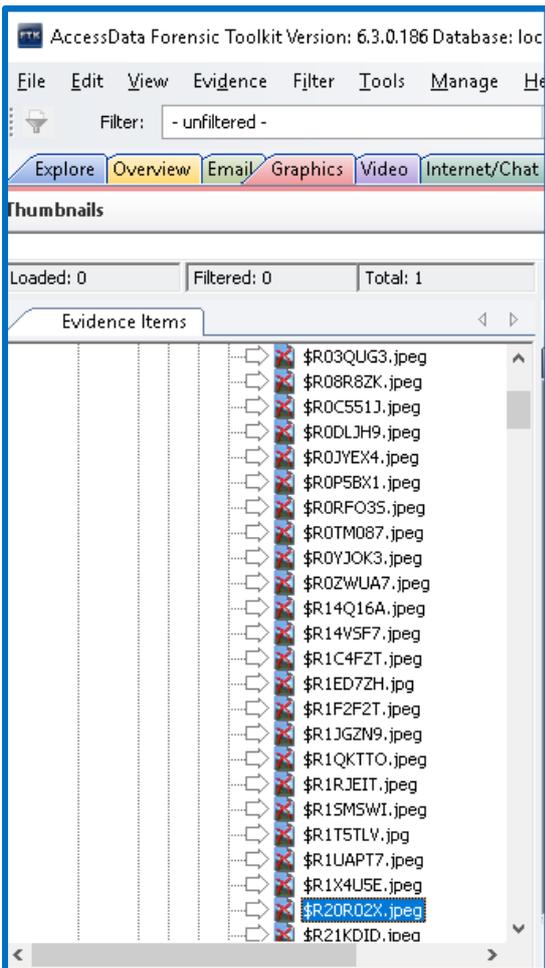


Figure 17: Kingston SSD- before FTK process.

When analysing the pictures, it was noticed that 6/25 of the files had data that had been replaced with zeros and 6/25 a different 6/25 displayed text instead of a picture. Some of the pictures appeared to have a Hex but no picture, this suggests that this picture may be recoverable because the Hex data is still present.

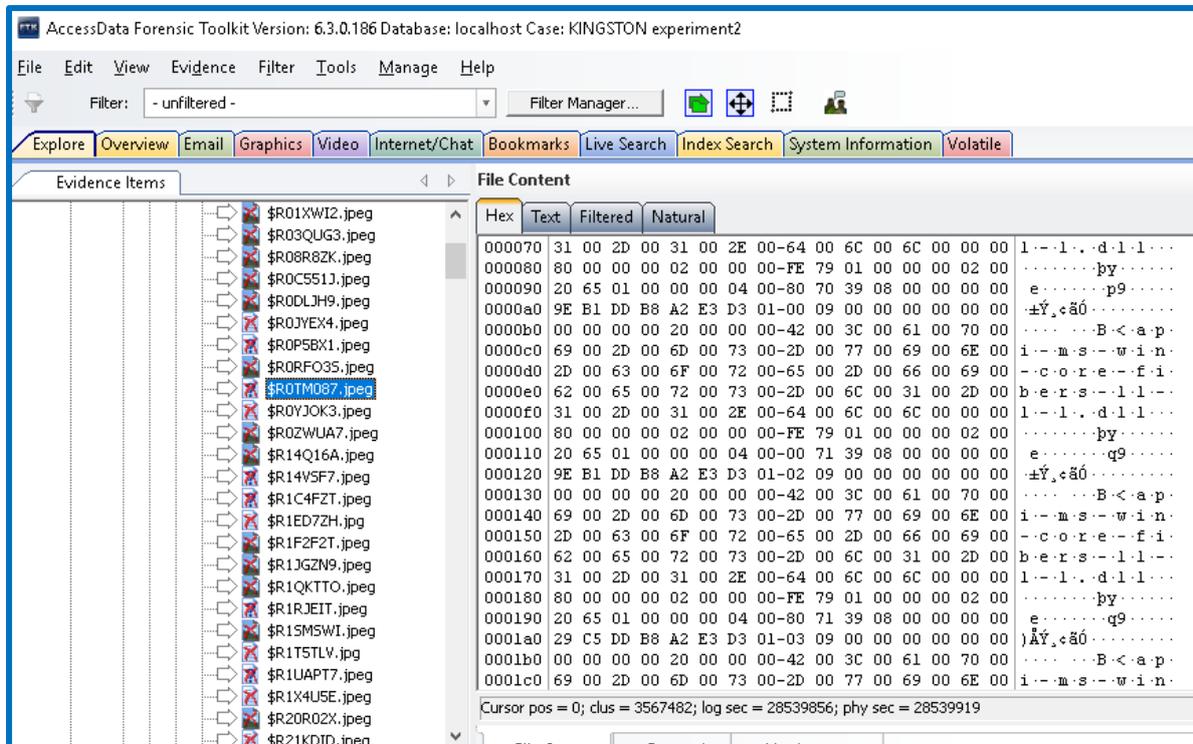


Figure 18: Kingston SSD- After FTK processed files

Figure 18 shows a file that has the Hex intact, but the picture is unable to be viewed.

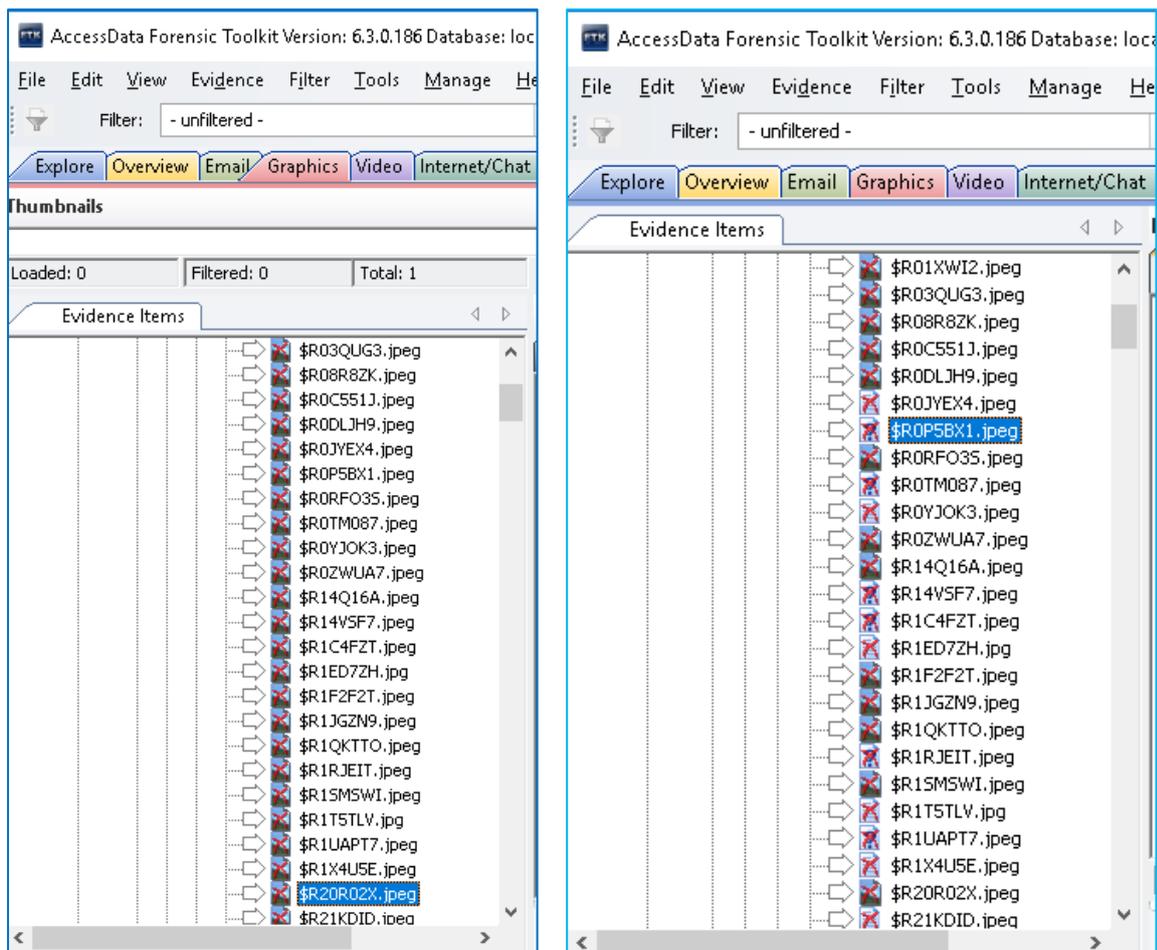


Figure 19: Kingston SSD- Side by side comparison

Figure19 This is what the files looked like after FTK had completed, in a side by side comparison.

Crucial SSD

The last SSD used in this experiment was the Crucial. It will be tested using the same methods as the previous two.

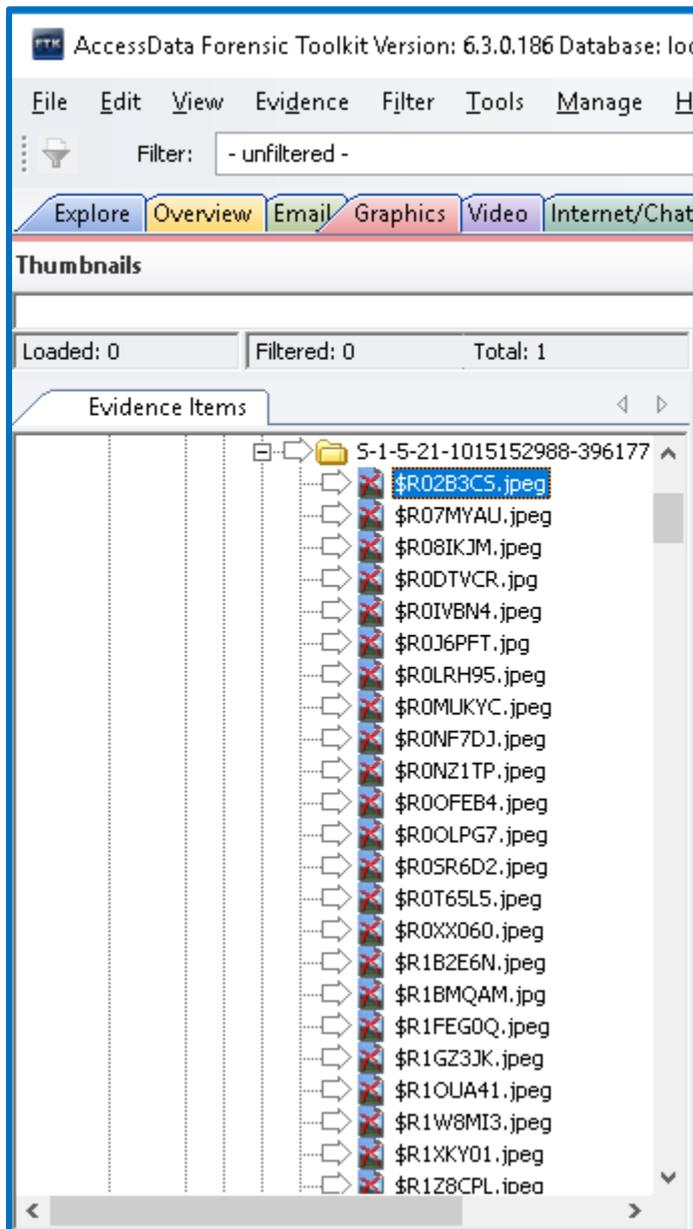


Figure 20: Crucial SSD- before images processed in FTK

Figure 20 shows what files looked like when first loaded into FTK. It displays the same issues involving the hex and zeros, as the previous two SSD tests.

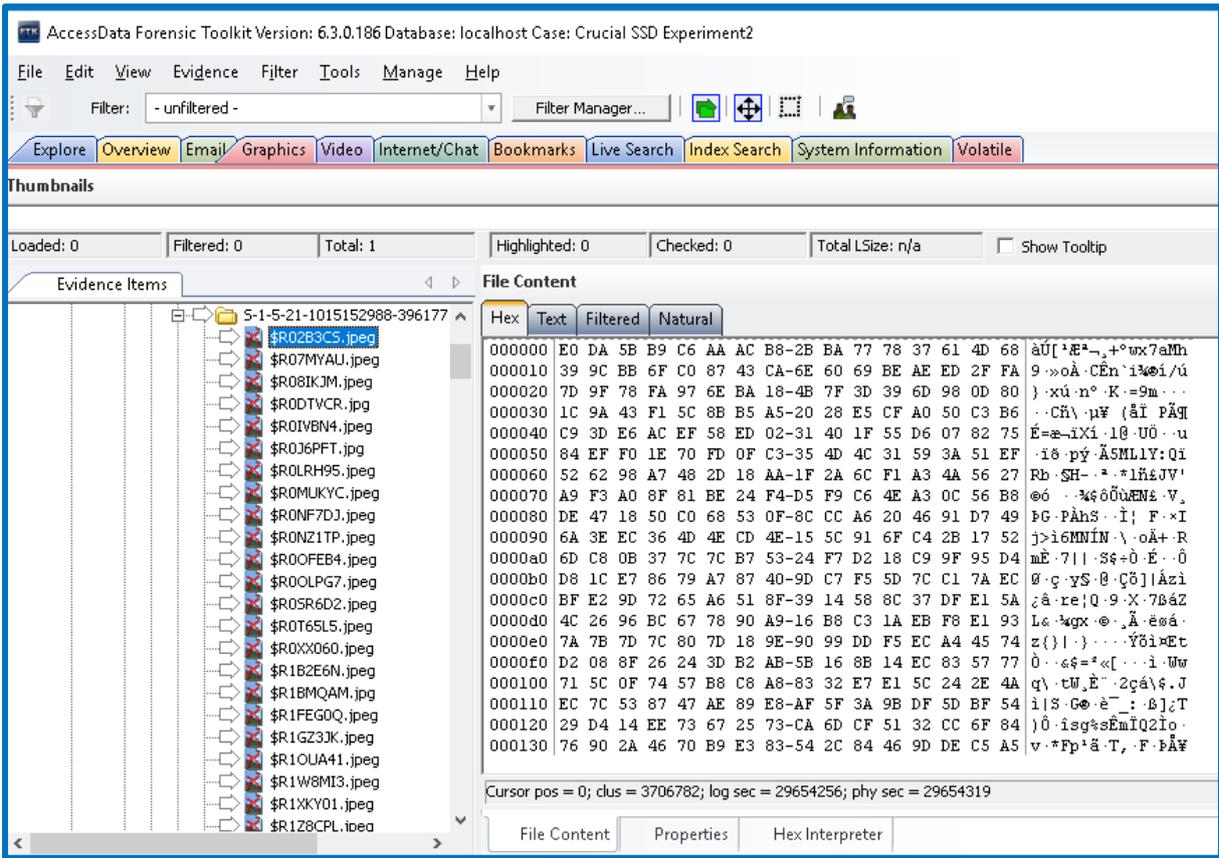


Figure 21: Crucial SSD- example of file in Hex format

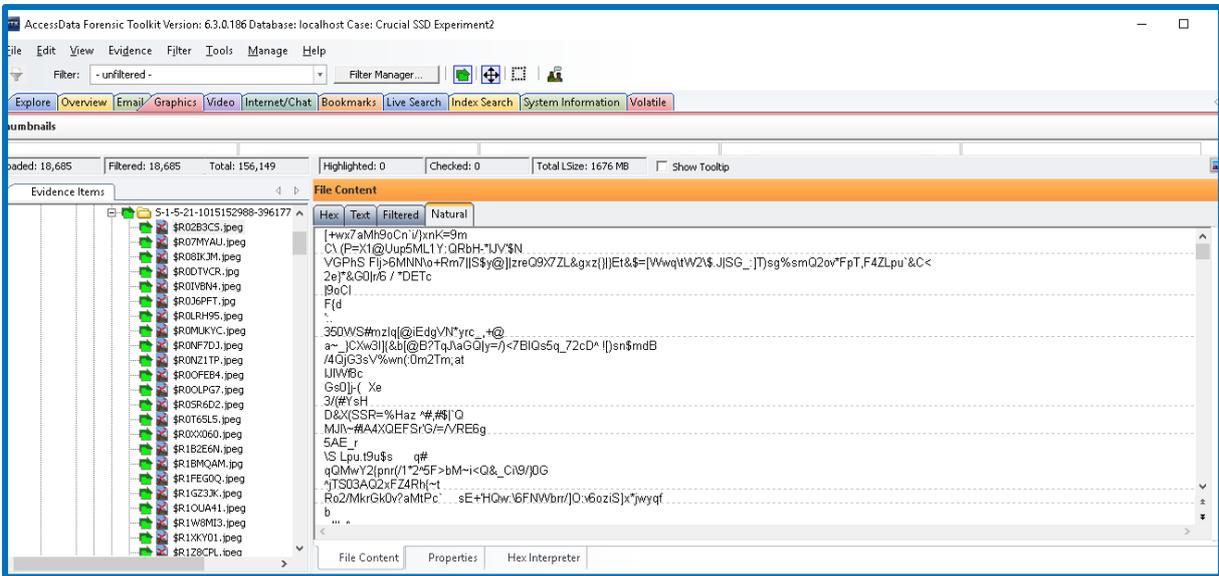


Figure 22: Crucial SSD- Example of file in natural format

Figures 21 and 22 show an example of the file that has hex data, but not the actual picture. This shows that the outcome for each SSD has shown the same issues throughout.

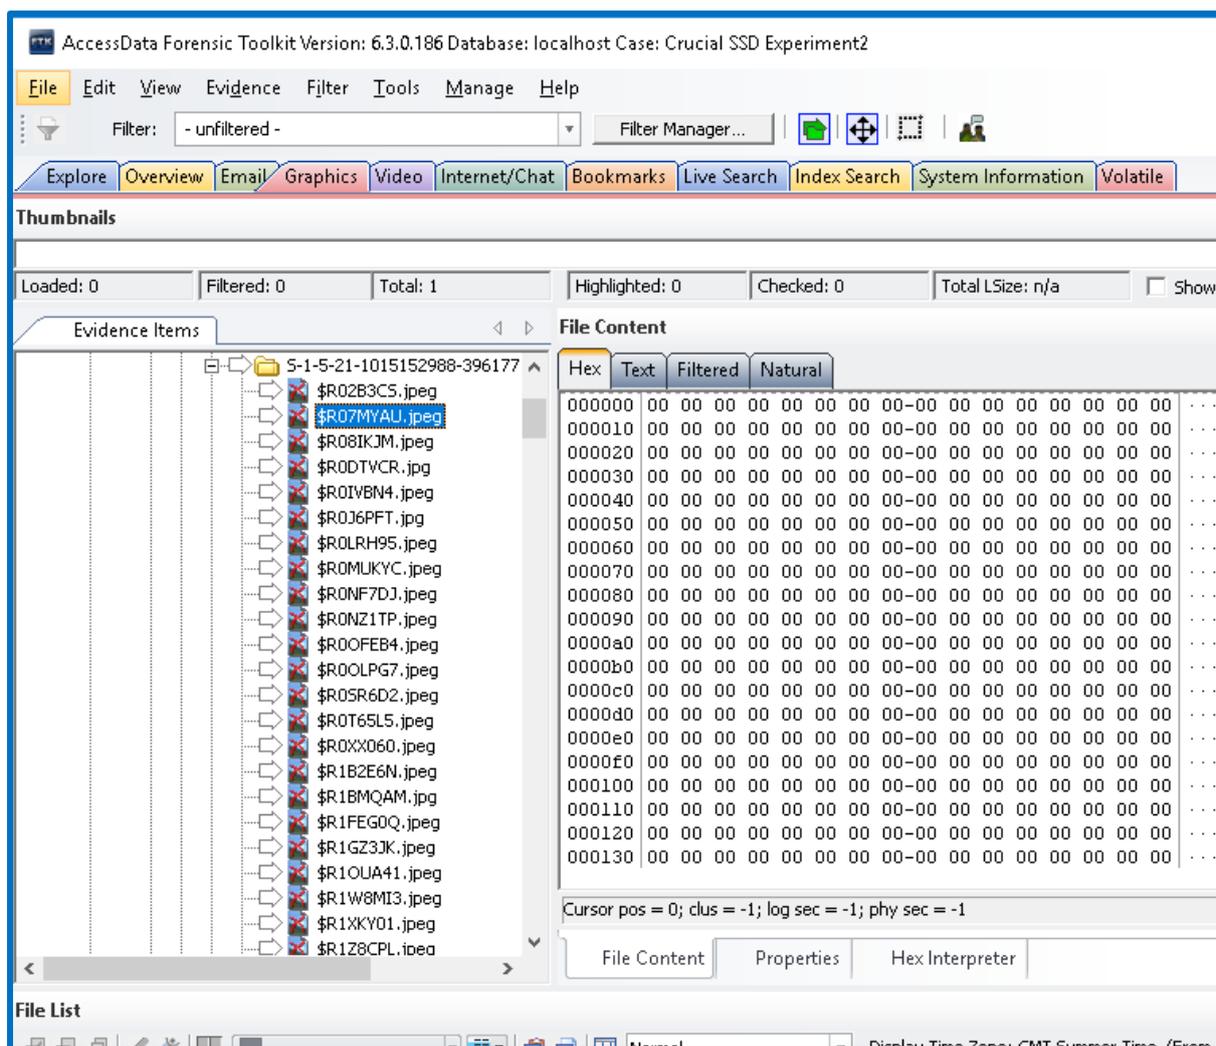


Figure 23: Crucial SSD- Example of file change after processing in FTK

Figure 23 shows an example of a file that has been zeroed out, but initially displayed as readable.

Actual results

SDD	How many files zeroed out	How many might be recoverable	How Many pictures can defiantly be recovered
Kingston SSD	143	41	127/311
Crucial SSD	117	32	162/311
WD Green	43	33	235/311

Table 3: Results of Experiment 2

Round up of Experiment 2

Out of the three SSDs tested, WD Green returned the largest number of recoverable files. This suggests that the garbage collection and TRIM function may run slower on this SSD, which is useful for forensic investigators, but in user terms, performs the worst.

This experiment revealed, many useful results that can be applied to solve some of the aims of this project. Firstly, it was noted that the TRIM type was definitely not DZAT, as the files were not entirely zeroed out as soon as the TRIM command was triggered. This meant that it could only either be DRAT or non-deterministic. However, it is difficult to tell which one it may be, but it is likely DRAT because it is more commonly used on modern SSDs such as the tested one.

An unusual result occurred when attempting to recover the data from all SSDs throughout the experiment. Initially the files appeared recoverable, however, upon closer inspection, some files were actually zeroed out unrecoverable. This shows that although it seems that self-corrosion is occurring on the SSDs, the data had already been through the garbage collection process and was just appearing to be in a recoverable state.

Similar to this, some files did not display the picture, but the hex was still intact. This further validates the fact that the TRIM type is either DRAT or non-deterministic. Unfortunately, there is a lack of literature available relating to how these types can be determined without information provided by the manufacturer.

6.1.3 Experiment 3

The main aim of experiment 3 was to address the secondary objective aim of this project and suggest a safe method of imaging an SSD.

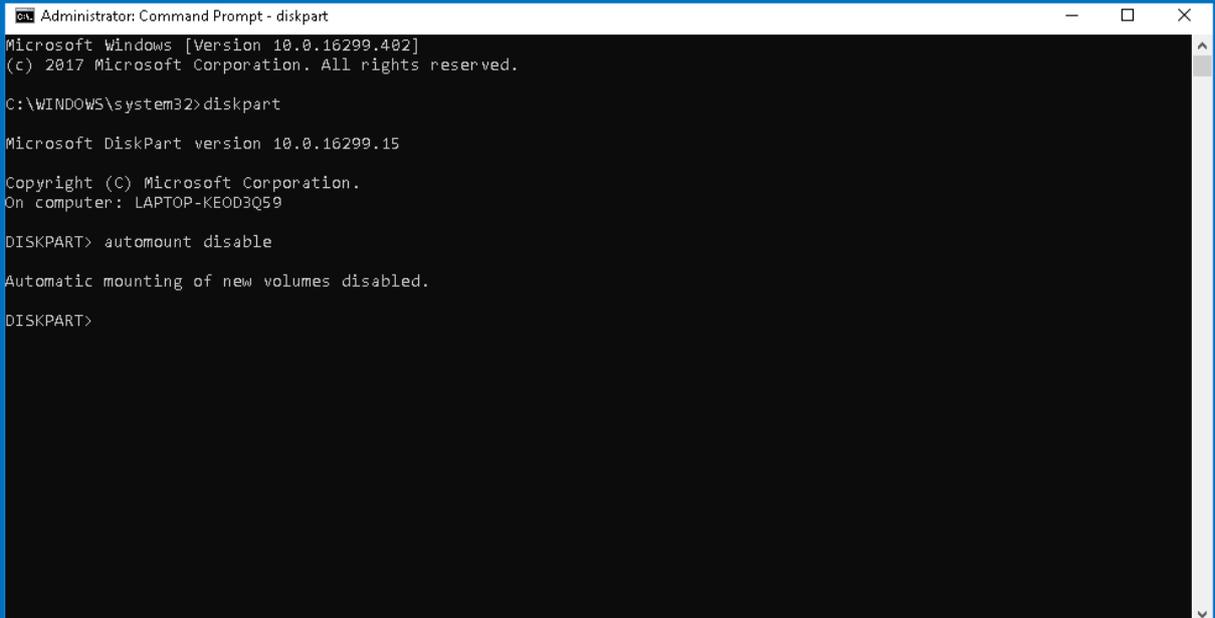
Research was conducted before experiments 1 and 2 to inform their methods. Experiment 3 has taken their outcomes into consideration to help focus and standardise the results. New research by Ferreira (2018) suggested that there was a potentially new solution for the issues face by modern digital forensics. It is based on the automount function which works by automatically inspecting devices that are connected to the computer. The study reveals that when automount is disabled, it has the potential to bypass the garbage collection function as the computer should not be registered to the device.

Although the study had positive results, it is notable that the SSD used was an older model. Within this project the modern M.2 SSDs are used, which have fast become the most popular type of SSD. This experiment will test the studies theory on the newer M.2 models to see if any impact can be made.

Experiment 3 Step by Step Methodology

1. An image will be created
2. MD5 hash the image
3. This image will be put on all SSDs using DD to ensure it is a bit for bit copy
4. 311 files will be deleted
5. SSD will immediately be removed
6. Automount will be switched off using diskpart command
7. SSD will be connected
8. Drive will be imaged
9. MD5 hash created
10. Image inspected for deleted files
11. Md5 hash checked.

Firstly, the automount was disabled by using the command line and inputting “diskpart”, then typing “automount disable”. This reduces the amount of interaction that Windows has with the SSD. However, the drive can still be seen by the FTK imager to be copied.



```
Administrator: Command Prompt - diskpart
Microsoft Windows [Version 10.0.16299.402]
(c) 2017 Microsoft Corporation. All rights reserved.

C:\WINDOWS\system32>diskpart

Microsoft DiskPart version 10.0.16299.15

Copyright (C) Microsoft Corporation.
On computer: LAPTOP-KE0D3Q59

DISKPART> automount disable

Automatic mounting of new volumes disabled.

DISKPART>
```

Figure 24: Diskpart in CMD

As this worked on older versions of SSDs it was expected to have an impact on the newer SSDs.

Expected results

SDD	MD5 Match	How Many pictures will be discovered
Kingston SSD	Yes	309/311
Crucial SSD	Yes	309/311
WD Green	Yes	309/311

Table 4: Expected results of experiment 3

Kingston SSD

The first SSD used for this experiment was the Kingston. This experiment is similar to experiment 1, apart from using automount to lessen the amount of interaction with Windows. The aim is to try and obtain better results than experiment 1 to compare and create a more efficient way of data recovery.

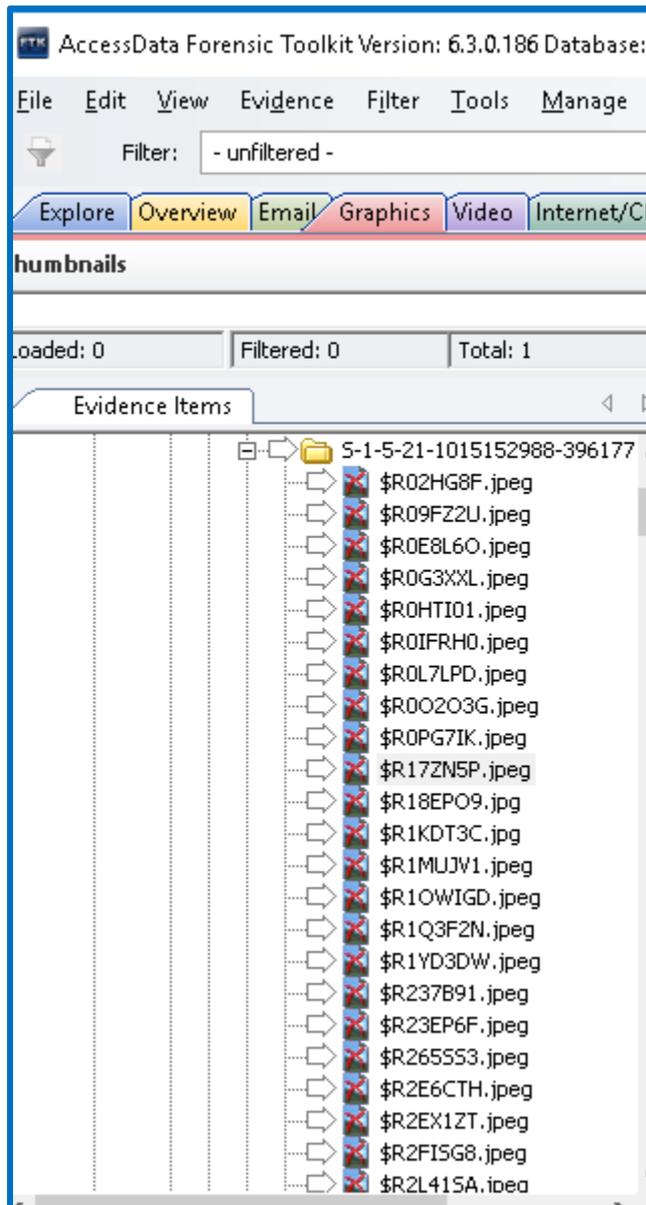


Figure 25: Kingston SSD- files when loaded in FTK

Figure 25 shows a screen shot of the files when first opened with FTK.

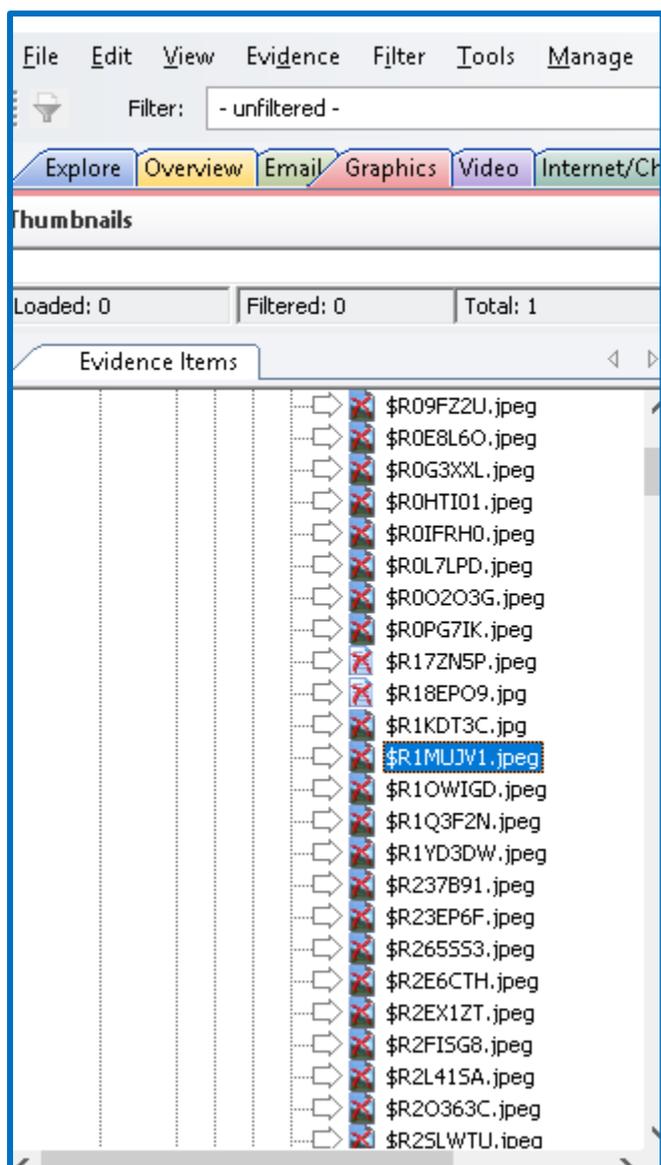


Figure 26: Kingston SSD- Files after being processed by FTK

Figure 26 shows a screen shot of after FTK had finished processing. This demonstrates that pictures became unrecoverable, the same as experiments 1 and 2. When these files were checked all had been filled with zeros.

However, there was a slight change within this experiment. Even files that displayed text when it was only supposed to display the picture were also filled with zeros. This means that these files are essentially unrecoverable, whereas if they retained the hex then there would be more chance of recovery.

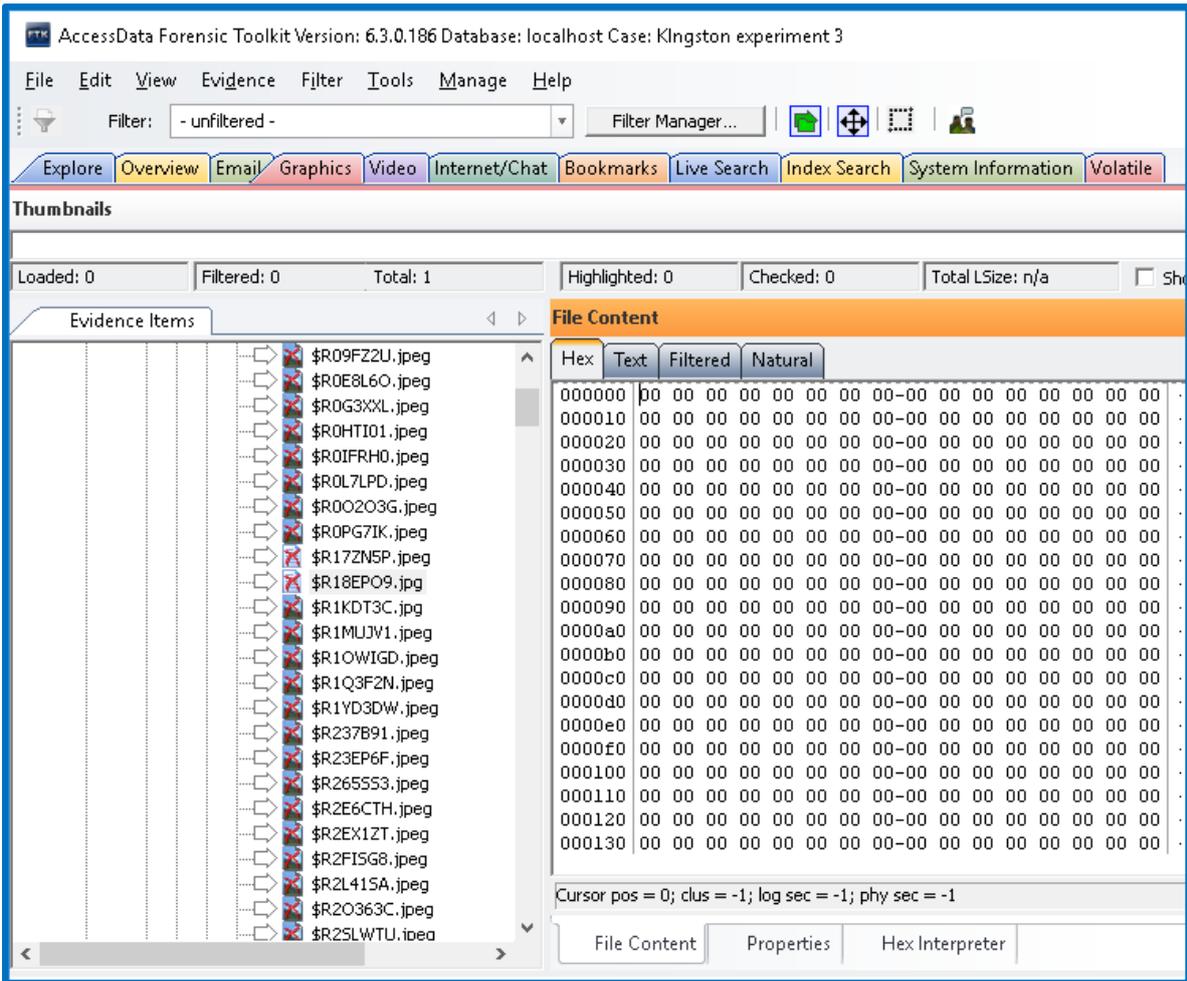


Figure 27: Kingston SSD- example of file with zeros

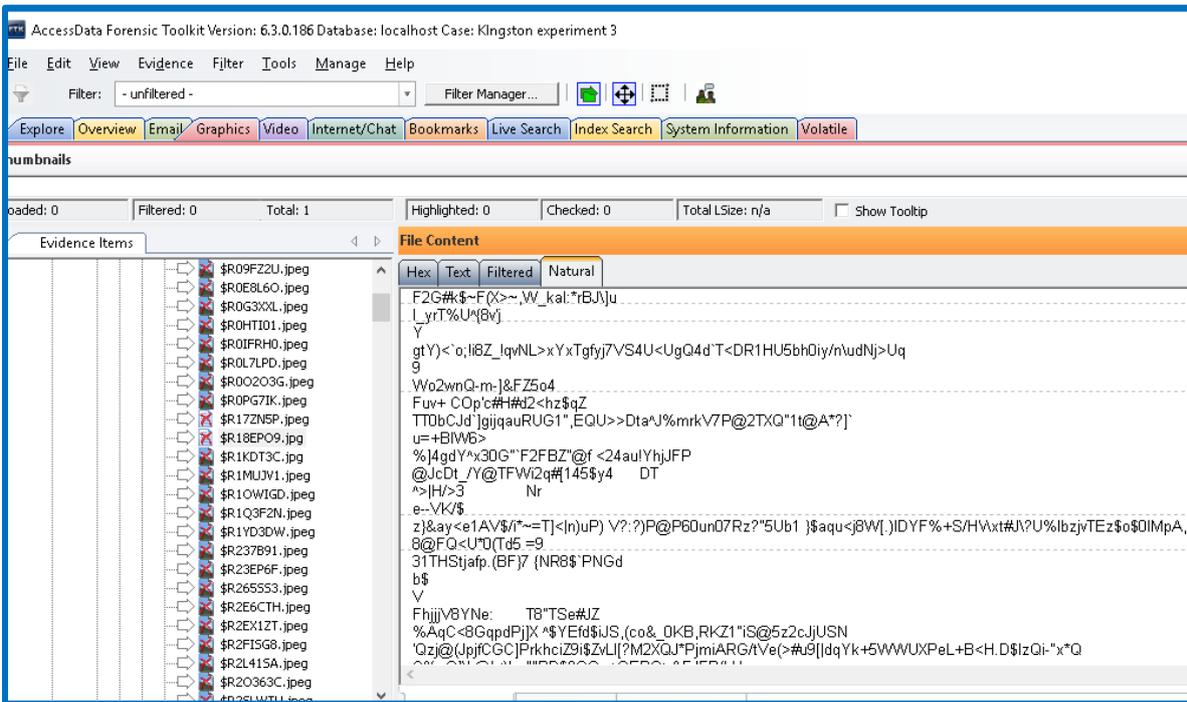


Figure 28 :Kingston SSD- File with its natural display.

WD Green SSD

The second SSD in the experiment is the WD Green SSD. This showed exactly the same issues as the previous SSD but also showed an increase of pictures recovered. Pictured in Figures 29 and 30 are the problems discovered on this SSD.

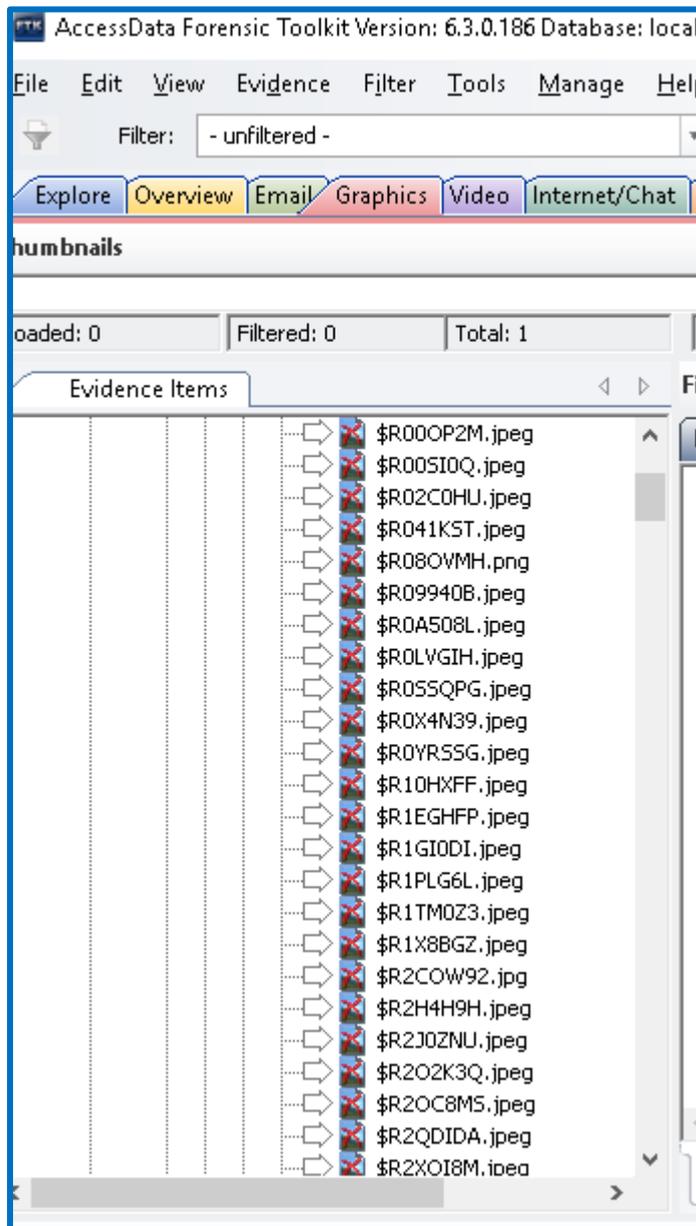


Figure 29: WD Green SSD- Files when first loaded into FTK

This picture shows what the files looked like when originally loaded into FTK.

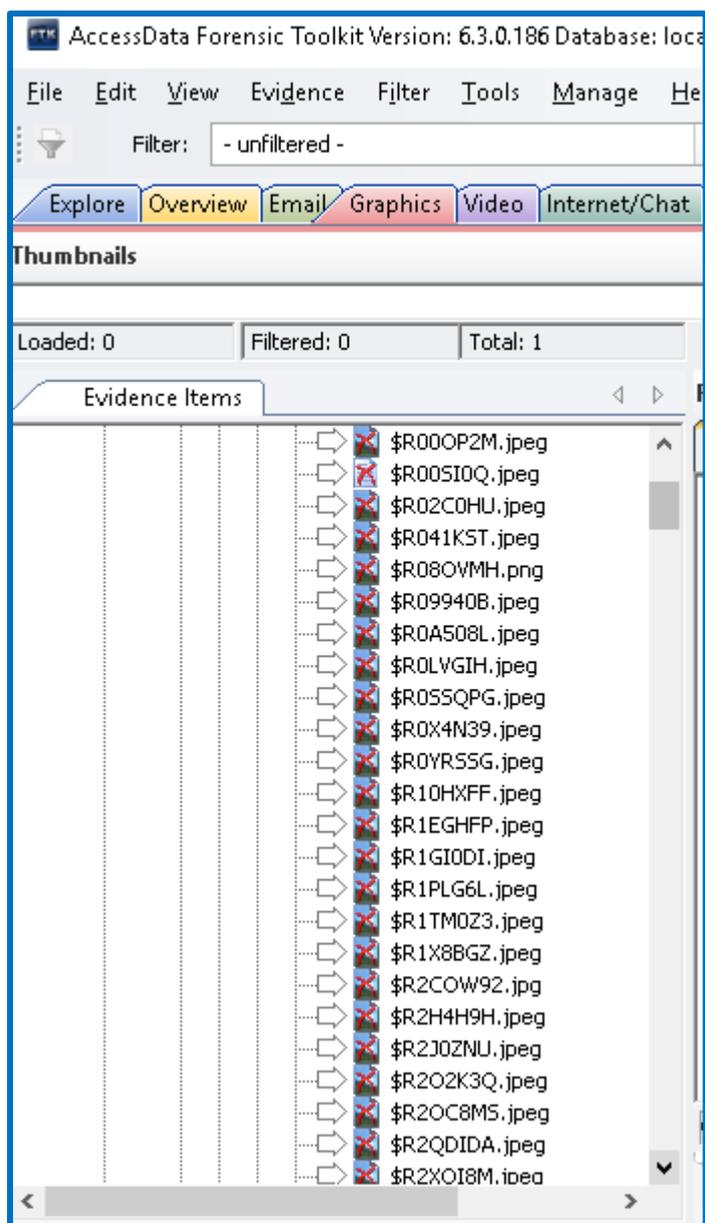


Figure 30: WD Green- Files after they were processed by FTK

Like the previous SSD the files changed and became unrecoverable once FTK had finished processing.

Unlike the previous SSD the WD green SSD didn't present zeroed out data when there was text displayed. This is exactly how it performed in experiment 1.

Crucial SSD

The last SSD in the experiment is the Crucial SSD. This SSD like the previous 2, show an increase in the amount of pictures that could be recovered. Pictured below, in Figures 31 and 32 are any issues that were found with the SSD.

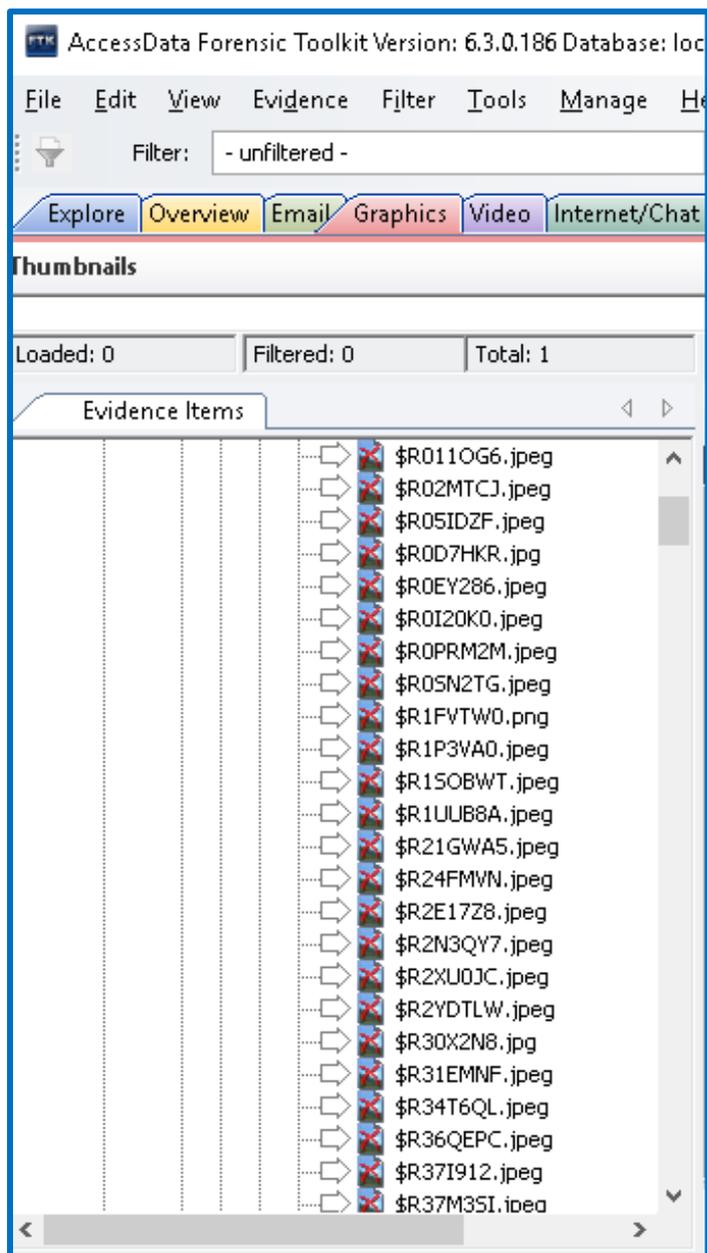


Figure 31: Crucial SSD-Files when loaded into FTK

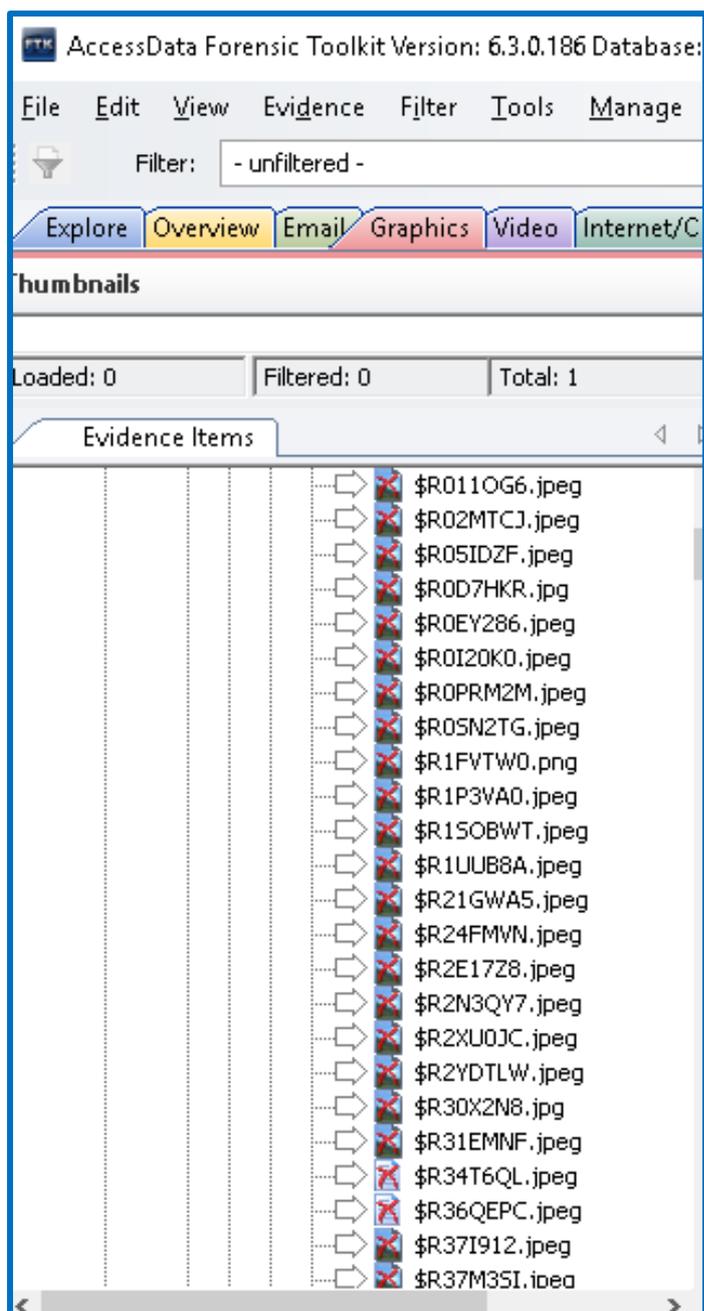


Figure 32: Crucial SSD- Files after being proceed by FTK

Actual results

SDD	How Many pictures will be recovered	Was it more than experiment 1
Kingston SSD	298/311	Yes
Crucial SSD	306/311	Yes
WD Green	303/311	Yes

Table 5: Actual results of experiment 3

Summary of experiment 3

After concluding experiment 3, it was clear that the amount of pictures that were recoverable increased whilst automount was disabled. This is a hugely positive development and one that can be taken further in future experiments. In comparison to the results of experiment 1, all of the SSDs held a higher amount of recoverable pictures. Reinforcing the validity of these results and providing more evidence to show that disabling automount could prove a useful technique to forensic investigations .

6.1.4 Challenges and Alterations Faced During Experimentation

There were multiple challenges that occurred throughout the experimental phase of this study. Although frustrating, it did provide insight into the complex works carried out by digital forensic investigators.

During the writing of the background, new research was published that offered inspiration for a complete rethink of experiment 3. The study by Ferreira (2018) described their theory for using auto-mount of the forensic workstation during SSD recovery to prevent TRIM and garbage collection from threatening the integrity of the evidence. However, this meant that experiment 3 was changed entirely to explore this theory of interest instead.

Initially, the image used in the experiments was to be provided by Cardiff University. However, two of the SSDs used in the tests, both 120GB, were too small to host this 110GB image. Therefore, an external hard drive of 80GB, was employed to create a bootable image on that was small enough to be used for all three SSDs in the tests.

Although the image was now usable for all SSDs, the imaging process itself took anywhere between 3 – 10 hours to complete. This became a serious issue as the allotted time for experimentation had to be extended significantly, especially since half way through the second set of experiments, the writeblocker stopped work and needed replacement. To ensure continuity throughout the project, it was important that the experiments were started from the beginning using the new writeblocker.

When the experiments were first designed, the original plan was to delete 30 pictures. However, once the experiments had started, it was noticed that 30 pictures were not enough to test the TRIM and garbage collection. Therefore, the number was increased to 311 which worked far better, but much time was lost before the change occurred.

Another original idea that needed to be changed, was the testing of SSDs at varying intervals to see which one performed best. After substantial amounts of research into this method, it was noted that other studies had not yielded any valuable results. However, it was decided that it would be tested anyway as there was potential for it to provide insight. After discovering it did not return any results of interest, it had wasted a considerable amount of experimentation time.

As discussed previously, there is a significant lack of information regarding the three different types of TRIM and which ones are used on certain platforms. Therefore, during the experiments, it was uncertain which TRIM type was at work. Since the results were not all

zeroed out, then the device couldn't have supported DZAT. Modern SSDs commonly have DZAT or DRAT, however, it is extremely difficult to tell the difference between non-deterministic TRIM and DRAT. This created difficulty in predicting and understanding some of the results obtained.

The final problem faced during experimentation was the unexpected space each SSD used up on the laptop. The laptop only contained 275GB and was expected to hold images for two 120GB SSDs and one 275GB SSD. Therefore, to allow for multiple imaging to be conducted an external hard drive of 1TB was used. However, towards the middle of experimentation, it was discovered that this would also prove too small to hold every image required. This meant that some images would have to be deleted to make room for new ones. This was not ideal, but the screenshots provided proof of work regardless.

7. Conclusion

At the beginning of this project, three primary aims and one secondary aim were outlined as deliverables of the research and experimentation. Due to a number of challenges faced during experimentation, these were adjusted accordingly to ensure the project had focus and results. The outcome and deliverable for each aim is as follows.

Primary

1. Determine the differences between the effects of each SSDs garbage collection and trimming processes on an image.

This method was altered from the initial plan, which stated that data would be recovered after varying intervals, instead recovering it immediately. This was done because when the intervals were included within the process, they did not provide any valuable information, therefore, they were not required to obtain results.

2. Understand what triggers garbage collection.

Research into the most current literature revealed that the workings of TRIM function triggered the garbage collection process the majority of the time. However, it was discovered that there are three different types of TRIM, each with varying effects on the recoverability of data after garbage collection has occurred. DZAT and DRAT and the most commonly used types, both are deterministic, meaning that the outcome of the deleted data is predictable. The third is the lesser used non-deterministic TRIM, which is more difficult to recognise as the outcome is less predictable.

However, in the experiments it was difficult to determine whether the type of TRIM on each SSD was non-deterministic or DRAT, due to the mixture of results returned after garbage collection. It is more likely to be DRAT as this is more commonly used in modern SSDs, but it cannot be guaranteed. This is important as the difference can have an effect on the recoverability of the deleted data.

3. Suggest ways of preventing garbage collection.

Experimentation revealed some interesting results when a diskpart command is used and automount is disabled. It appeared as though more files were recoverable. This could be due to the lack of interaction between Windows and the SSD because the automount is disabled.

Although there was no solid evidence for the writeblocker preventing garbage collection, there were suggestions in the results that showed the SSDs self-corroded faster without the presence of the writeblocker. For the first experiment, with the writeblocker connected, the Kingston SSD showed that 291/311 files were recoverable and without the writeblocker connected, on the second experiment only 127/311. Further testing on this would be needed to validate these results.

Secondary

1. Suggest safe methods of imaging.

Results from experiment 3 show that it may be possible to limit the damage of garbage collection but disabling automount. This could be incorporated into their practiced forensic investigation techniques to help them to recover as much data as possible.

The equipment required:

- M.2 to SATA SSD Enclosure x1
- Tableau Forensic PCIe Bridge x1
- FTK Imager
- Host computer

The step by step method:

All these steps are under the assumption that the SSD M.2 has been removed from the computer /Laptop.

1. Turn on Host computer
2. Open CMD as administrator
3. Type "diskpart" command and wait for new window to open
4. Type "automount disable"
5. Restart computer
6. Open FTK imager and fill in file name details
7. Connect M.2 Sata SSD to the Tableau forensic bridge
8. Insert M.2 SSD to M.2 to SSD enclosure
9. Turn on Tableau forensic bridge
10. Click on the start button on FTK

Automount will be switched off using diskpart command.

With reference to the experiment results a number of suggestions and recommendations have been realised. The police force and the forensic community currently struggle to know when garbage collection and TRIM functions have been initiated. The multiple manufacturers further this difficulty by creating SSDs that have different triggers for these functions and do not disclose any technical information to the public, professionals or even the police force. Therefore, there seems to be a clear requirement for conversations between the manufactures and the police force or the forensic recovery software companies to come up with solutions.

By liaising with forensic recovery software companies, such as FTK, there is a possibility that commands to delay garbage collection, immediately suspend garbage collection or to read data that has been trimmed, can be created. This way the manufacturers can maintain secrecy for copyright interests, but still contribute towards finding solutions to the problems.

It seems nonsensical for researchers to be desperately trying to find solutions to the issues of data loss when it is highly likely that the companies that create the varying models of SSDs,

already hold the answers to. The majority of problems currently faced by local police forces investigating criminal behaviour, or forensic investigators recovering valuable data, can be solved by simply starting a confidential conversation with manufactures. Although this project revealed some interesting aspects of garbage collection and TRIM function, it is only just scratching the surface and if the manufacturers collaborate with forensic teams then many more unknowns can be revealed.

8. Future Work

Through the research and experimentation conducted within this project, it was discovered that there are many shortcomings relating to the current understanding of forensic investigation of modern SSDs. This is particularly notable when studying the current ACCPO guidelines used by the UK police force to inform forensic investigators. The ACCPO document states the procedures followed by digital forensic investigators to carry out a formal investigation on hard drives and how to present the evidence in a state that is valid for use in court.

Although this guidance offers a good template for digital forensic recovery in general, it is far too vague to provide effective instructions on how recover deleted data from SSDs. As explained in the research, SSDs are rapidly developing and the manufacturing companies are becoming more secretive with the design specifications. This has meant that forensic investigators and police forces have not been able to maintain their understanding of SSD functions, such as garbage collection and TRIM, that affect data loss.

Therefore, it is recommended that using the findings of this paper, alongside further research, ACCPO should be updated to offer a set of standard guidelines for each type of SSD. However, it is realised that for this solution to arise, the manufacturers must be willing to work alongside forensic specialists. This would encourage communication between the two, allowing forensic investigators access to important information that could make a significant difference in a prosecution case.

This subject would benefit from further research into the more complex operations of garbage collection. This project was based upon the deletion of the contents of one folder on a drive, however, it would be interesting to conduct an experiment involving a drive with multiple folders and files saved on various blocks. In addition to this, files could be deleted in various amount and from different places, as opposed to all at the same time from one location. This would be helpful to assess how the garbage collection occurs in a scenario even closer to real life.

9. Reflection

Upon reflection of the initial project plan, it was clear that the tasks set were too vague and slightly ambitious for the time and resources available within this project. The initial plan was based on information that Gwent Police had provided relating to the problems that they faced whilst attempting to recovery data from SSDs. It wasn't until the research phase of this project, that the extent and variations of the issues were realised. However, the research allowed the project to gain focus and direction, leading to the development of experiments that were useful but tangible within the given time frame.

This realisation altered the expected time frames of the experiments according to the Gantt chart. The first three weeks of research remained the same, but the experiments began later than planned to ensure that they were fully researched and worthwhile.

The write up for this paper remained on track as there was a substantial amount of background research to be completed. Because the subject broached within this project is so broad and follows a long history of relevant information the paper contains a considerable amount of general information about SSDs and their functionality. Although it is all relevant to the issues posed, collecting and writing the information took up a significant amount of time that could have been used experimenting. However, because of the extensive research effort that was involved in creating this project, it has provided an excellent basis for future work to be conducted as much of the relevant and up to date information is collated all within this one paper.

The results and evaluation write up took slightly longer than expected as the continuing changes to the experiments were unforeseen at the time of creating the Gantt chart. However, overall the timeline of work during this project was closer to the predicted Gantt chart than expected. In addition to this, the project has provided Gwent Police with some potential solutions for the questions that were asked and has identified a safer way to investigate SSDs.

Because the subject of SSD forensics is so new, studies into the complexities and prevention of self-corrosion are continually being conducted. Even during this project, a number of research papers proposing new ideas and methods for digital forensic investigation of SSDs were being published. Therefore, it is likely that there will be continual development in this area, offering more suggestions and experiments to be conducted.

10. Glossary

- ACPO – Association of Chief Police Officers, guidelines for best practice.
- Automount – When Windows automatically registers any device when connected to a computer.
- Blocks – Made up of a collection of pages.
- CMD – Command prompts.
- Controller – Embedded processor that connects the NAND flash memory and the computer.
- DaP – Deconstruct and preserve.
- DRAT – Type of deterministic TRIM.
- DZAT – Type of deterministic TRIM.
- FTK – Forensic toolkit, used for recovering data.
- Garbage collection – A function on an SSD that permanently deletes data marked for deletion by the computer.
- HDD – Hard disk drive.
- Image – A copy of a drive.
- M.2 SSD – SSDs created after 2015.
- MD5 Hash – An algorithm that creates a has value, used to ensure integrity of data.
- NAND flash memory – Type of non-volatile storage technology.
- Non-deterministic – Type of non-deterministic TRIM.
- Pages – The smallest unit of an SSD.
- Raw image format/DD – Bit by bit copy of the original.
- SATA – Serial ATA, a computer bus interface.
- Self-corrosion – Even without computer instructions, the SSD will permanently delete data from the drive.
- Slack space – The space left over when smaller data overwrites larger data.
- SSD – Solid state drive.
- TRIM – A command sent from the computer to the controller which locates the data marked for deletion.
- Writeblocker/forensic bridge – Device that gains read-only access to storage devices.
- Zeroed out – Where data only appears as zeros.

11. References

- Bell, G. and Boddington, R. (2010). Solid State Drives: The Beginning of the End for Current Practice in Digital Forensic Recovery?. *Journal of Digital Forensics, Security and Law*.
- Computerhistory.org. (2018). 1990: Magnetoresistive read-head HDD introduced | The Storage Engine | Computer History Museum. [online] Available at: <http://www.computerhistory.org/storageengine/magnetoresistive-read-head-hdd-introduced/> [Accessed 6 Mar. 2018].
- Digital-detective.net. (2018). [online] Available at: http://www.digital-detective.net/digital-forensics-documents/ACPO_Good_Practice_Guide_for_Digital_Evidence_v5.pdf [Accessed 11 May 2018].
- Ferreira, J. (2018). Forensic Acquisition Of Solid State Drives With Open Source Tools. *Forensic Focus*. [online] Available at: <https://articles.forensicfocus.com/2018/03/13/forensic-acquisition-of-solid-state-drives-with-open-source-tools/> [Accessed 3 May 2018].
- Foote, K. (2017). A Brief History of Data Storage. *Dataversity*.
- Garfinkel, S. (2010). Digital forensics research: The next 10 years. *Digital Investigation*, 7, pp.S64-S73.
- Goda, K. and Kitsuregawa, M. (2012). The History of Storage Systems. *Proceedings of the IEEE*, 100(Special Centennial Issue), pp.1433-1440.
- Gubanov, Y. and Afonin, O. (2012). Why SSD Drives Destroy Court Evidence, and What Can Be Done About It. *Forensic Focus*.
- Gubanov, Y. and Afonin, O. (2016). SSD and eMMC Forensics 2016 - Belkasoft. *Forensic Focus*.
- Hgst.com. (2018). DeskStar 7k1000. [online] Available at: https://www.hgst.com/sites/default/files/resources/Deskstar7K1000_010307_final.pdf [Accessed 7 Apr. 2018].
- Mitchell I., Anandaraja T., Hara S., Hadzhinenov G., Neilson D. (2016) Deconstruct and Preserve (DaP): A Method for the Preservation of Digital Evidence on Solid State Drives (SSD). In: Jahankhani H. et al. (eds) Global Security, Safety and Sustainability - The Security Challenges of the Connected World. ICGS3 2017. Communications in Computer and Information Science, vol 630. Springer, Cham
- Pixabay.com. (2018). *Images Images · Pixabay · Download Free Pictures*. [online] Available at: https://pixabay.com/en/photos/?q=images&hp=&image_type=all&order=&cat=&min_width=&min_height= [Accessed 10 May 2018].
- Rajan, R. (2012). Evolution of Cloud Storage as Cloud Computing Infrastructure Service. *IOSR Journal of Computer Engineering*, 1(1), pp.38-45.

Stevens, L. (1981). The Evolution of Magnetic Storage. *IBM Journal of Research and Development*, 25(5), pp.663-676.

Thatcher, J. (2009). NAND Flash Solid State Storage Performance and Capability – an In-depth Look. *SNIA*.

Wilkes, M. (1980). Computers into the 1980s. *Electronics and Power*, 26(1), p.67.

Www-03.ibm.com. (2018). *IBM Archives: IBM 1311 disk storage drive*. [online] Available at: https://www-03.ibm.com/ibm/history/exhibits/storage/storage_1311.html [Accessed 3 Mar. 2018].