ESTABLISHING PROFESSIONAL GUIDELINES FOR SSD FORENSICS: A CASE STUDY

JAY JUNICHIRO UCHIYAMA (B.Bus)

A thesis submitted to the graduate Faculty of Design and Creative Technologies
Auckland University of Technology
in partial fulfilment of the
requirements for the degree of
Master of Forensic Information Technology

School of Computer and Mathematical Sciences

Auckland, New Zealand
2014

Declaration

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which, to a substantial extent, has been accepted for the qualification of any other degree or diploma of a University or other institution of higher learning, except where due acknowledgement is made in the acknowledgements.

.....

Jay Junichiro Uchiyama

Acknowledgements

This thesis was carried out during the year of 2013 at the Faculty of Design and Creative Technologies in the School of Computing and Mathematical Sciences at Auckland University of Technology, New Zealand. Support was received from many people throughout my postgraduate education.

It is a pleasure to thank those who made this thesis possible, I would like to acknowledge the significant support and encouragement provided by my family during the course of conducting this research as well as throughout my entire postgraduate study.

I owe my deepest gratitude to my supervisor, Prof. Brian Cusack. Without Dr. Cusack's continuous optimism concerning this project, inspiration, and guidance this thesis would hardly have been completed.

I am deeply grateful to my employer PricewaterhouseCoopers and Forensic Services Director, Campbell Mckenzie for always being a motivational influence in all aspects of my study and research. Mr Mckenzie has shared advanced knowledge and commercial insights of digital forensic procedures, and allowed my access to numerous forensic tools.

Lastly, I am indebted to many of my colleagues, especially Alain Homewood, Roman Ammann, Craig Calderwood, Shubham Sharma, New Zealand Police Electronic Crime Laboratory Digital Forensic Analysts Mark Simms and Ben Knight who have been supportive and provided inspiration in all manners for my chosen field of study. Also thanks to the proof reader whom has profoundly improved the composition of this thesis.

Abstract

The aim of this thesis is to investigate and examine the present status of solid state drive (SSD) forensics, and to establish a professional guideline for forensic investigators who are required to preserve and recover data stored on SSD in a forensically acceptable manner.

In the first part, results of a literature review of computer storage devices, data recovery methods, and forensic guidelines were presented. The literature review determined how SSD is architecturally different from a magnetic hard disk drive (HDD), but existing forensic guidelines and procedures were developed based mainly on HDD technology. SSD is widely accepted by consumers but not well integrated into the forensic guidelines, despite several automated evidence-destruction functions, which were embedded for performance enhancement purposes, have been explicitly discussed by forensic and data recovery experts.

The thesis then identifies the gaps amongst well repute forensic guidelines and further outlines the structure of a compound guideline which recognises issues raised by SSD to maximise the chance of data recovery. Specific processes were identified and data recovery rate was measured for testing.

In conclusion, the thesis argues that existing forensic techniques and guidelines are incapable of suppressing the SSD's self-destructive behaviour and alternative method of SSD data preservation must be developed.

Table of Contents

Chapter 3 – Research Methodology

Chapter 4 – Research Results

Chapter 5 – Discussion

Chapter 6 – Conclusion

List of Tables

List of Figures

Chapter 1

Introduction

1.0 ESTABLISHING THE IMPORTANCE OF THE TOPIC

In recent years, there has been an increasing interest in solid state drive (SSD) forensics. SSD core technology has been around since 1989, but it is fairly recent that the technology became stable and affordable to generic consumers. The price is still higher than traditional hard disk drive (HDD), but lower power consumption and incredibly fast response times attracts many keen users across the globe. Typically HDD is being replaced by SSD as a primal computer storage device for the higher performance gains.

A global trend in primary computer storage devices is shifting from HDD to SSD has become a major threat to computer forensic investigations. Savvy consumers are generally aware that unlike HDD, SSD suffers from wearing and there is a limit to the number of writes can be performed. The write endurance limit issue has been widely discussed because such limitation never existed with magnetic storage devices. Manufacturers implemented a built-in process known as wear-levelling to minimise excess usage of certain parts of SSD and evenly balance out the usage, and users became less concerning of the write limit. However, this is causing catastrophic damage to the residual data.

Traditionally a duplicate image of seized HDD is made as part of the evidence preservation process. The image is known as a forensic copy and mathematically calculated hash values are used to verify the copy and the original are the exact duplicate. Any forensic analysis is performed on the forensic copy. The exact duplicate even allows recovery of deleted data without risking alteration to the original HDD. The correct procedures are carefully documented and made publically available as computer forensic guidelines or the best practice through number of reputable academic and government institutions.

There are two significant issues immediately raised when SSD was tested. Firstly wear-levelling and other optimisation processes are occurring without any user intervention. Automation of these processes means that there is no way of stopping such process at a user's will. The optimisation processes remain active

while the forensic duplication is proceeding. From a data duplication point of view, the state of original data became dynamic on SSD while HDD was static. The automated background process also gave dynamic mathematical hash value calculations on SSD with the result that a copy will never be able to verify the completeness.

The obsolete area where the residual of deleted data remains is purged by the optimisation processes. Any data recovery tools and techniques rely on residual data for carving and reconstruction of deleted files. The same method is applicable to computer forensics where the forensic image is designed to copy everything regardless of its state. The automated processes are purging traces of information potentially critical to evidence. Worst of all, these processes are active as soon as power is supplied to the device, and there is no way of stopping them.

1.1 HIGHLIGHTING A PROBLEM IN THE FIELD OF STUDY

Several studies have produced research on SSD forensics legal challenges, and difficulties of recovering data, but no thorough solution is made available. SSD forensics is a phrase used to describe any forensic tools and techniques applied to process data stored on a SSD.

Bell and Boddington (2010) reported that existing acquisition methods are no longer valid for SSD devices and they are unable to function in many situations. Their experiment demonstrated SSD have the capacity to destroy evidence under their own volition unless specific instructions are sent from a computer. These authors are not the only researchers raising the concern regarding the garbage collection functionality (self-destructive behaviour) that manufacturers embedded for performance optimization.

Bednar and Katos (2011) indicated that due to the self-destructive behaviour, procedures defined in the Association of Chief Police Officers' (ACPO) digital forensic guide book no longer applicable for handling SSD devices. Instead, if a SSD was found powered "on", authors suggest that a SSD to be immediately unplugged to preserve the evidence.

In contrast to Bednar and Katos's study, Freeman and Woodward (2009) researched on secure deletion on a SSD demonstrated that SSD's wear-levelling

function provides efficient functionality and speed, but it is also an obstacle to ensuring the removal of all information from the drive.

King and Vidas (2011) published their research for SSD data retention analysis and discuss the data recovery problem faced by forensic examiners due to the ATA8 TRIM command. The experiment shows that with TRIM enabled, only up to 27% of blocks were recoverable but without TRIM (such as Windows XP operating system) nearly all data is recoverable. It is interesting that King and Vidas mention that Bell and Boddington's experiment used small (10KB) sample file sizes and that reduced the usefulness of the results.

Prior to SSD, Stokes (2008) presented how data can be recovered from mobile device NAND flash. Breeuwsma et al (2007) demonstrated data recovery on USB memory sticks. Stokes and Breeuwsma both used JTAG (direct memory access) acquisition techniques, which requires extensive modifications to the original hardware and the procedures are irreversible. An approach such as JTAG may allow physical access to the data, but raises numerous questions in regards to forensically acceptable practice.

SSD forensics is a current topic in the field of computer forensics. Questions have been raised about the safety of prolonged use of existing forensic tools and techniques developed for magnetic storage devices over decades. Further research is required to establish in-depth knowledge of best practice for SSD forensics.

1.2 RESEARCH FOCUS AND OBJECTIVE

The objectives of this research are to test various acquisition tools and techniques on SSD, determine the possibility of improving data recovery on SSD in forensically acceptable manner, and develop guidance for best practice.

Literature for similar studies will be reviewed to identify research gaps and potential issues with current guidelines. There are number of guidelines commonly referenced across the world, which will be identified and critically assessed to test SSD forensic capability, suitability, and readiness.

As a part of research design, a new guideline process flow will be drafted, and research experiments will be designed to test if the revised and compound guideline can help improve the rate of recovery and resolve issues SSD has brought to digital forensics. Research questions and associating hypotheses will

be derived from the literature review and presented with the research methodology.

The experiment is to use multiple sample SSDs as well as one HDD. The result from HDD will be used as a control. Results from the literature review will be used as benchmarks and the test results from this study will be compared. Improvement in data recovery will be measured by quantified rate of successfully recovering data. Successfully recovered data means the data is loadable and retains evidential value. Criteria will be set for data collection, which will be analysed then discussed.

Quantitative analysis will be used to measure the improvement in data recovery. The results are analysed and presented in tables, figures, and charts. The qualitative analysis will be then evaluated to derive logical reasoning for the results and the observations will also be stated. Defined hypotheses will be tested and research questions will be answered as part of the discussion section.

1.3 THE THESIS STRUCTURE

The Thesis is divided into six parts. Chapter 1 has introduced background information to promote the context and the significance of the topic selected for this research. It is followed by acknowledgement of the relevant studies that show the focus and the scope of research field. The objectives were set to clarify the research statement, and the scope of the work explained.

Chapter 2 will present the literature review. The purpose of the literature review is to constitute clarity, relevance, refinement, and identification for this research. A detailed review of both HDD and SSD storage devices from raw materials to software architectures will be discussed as a foundation for this research. Data recovery techniques are explored and reviewed as relevant to the field of digital forensics. In addition existing forensic guidelines from various well reputed institutions are reviewed and re-organised into process level flow-charts for capability gap identification.

Further reviews will be continued in Chapter 3 to establish how others have done similar research. Relevant studies are reviewed to identify suitable research models and methods. The research plan including research questions, hypotheses, and procedures will be addressed in Chapter 3. Details of test

requirements, methodologies, scope, limitation, and analytical methods will also be defined.

Chapter 4 will provide actual observations of the results including statistics, tables, and graphs, while Chapter 5 discusses interpretations and statements about the meaning or significance of the observations. Answers to the research questions and hypotheses testing are made in Chapter 5.

Finally Chapter 6 will present the conclusion and recommendations for further research. A complete list of references follows.

Chapter 2

Literature Review

2.0 INTRODUCTION

Digital forensics is a process where hypotheses are developed and tested for an event which involves electronic information called *digital evidence*. Carrier (2011) defines digital forensic investigation is a "process that uses science and technology to analyse digital objects and that develops and tests theories, which can be entered into a court of law (admissibility), to answer questions about events that occurred". A typical forensic process consists of four main phases of collection, examination, analysis, and reporting (Kent, Checalier, Grance, & Dang, 2006).

The United States National Institute of Standards and Technology (NIST) describes that data collection involves two stages, those being source data identification and data acquisition. Data or digital evidence must be seized in a forensically acceptable manner. The term "forensically acceptable" means "admissible in court" in this context. Therefore vigorous preparation work to identify evidence source at the crime scene is critically important.

The acquisition stage is also known as a "*preservation of evidence"*, and involves three-step processes. These are "prioritisation", "acquisition", and "verification" (Kent, Checalier, Grance, & Dang, 2006). Source of evidences are prioritised based on their likely value, volatility, and the amount of effort required. Digital forensic methodologies, guidelines, and best practices are commonly implemented for a data acquisition plan. The best available forensic tools are utilised, and a duplicate copy known as a *"forensic image"* is created. This format allows forensic analysts to work on the data but retain and validate data integrity through the forensic investigation lifecycle.

The storage device market is facing a major transition phase as the types of storage change. HDD has been the mainstream data storage device for both the consumers and the enterprise users since 1980's. Advent of random access memory (RAM) technology is beginning to replace the domination of HDD in the high-end storage sector with NAND flash based storage devices known as Solid

State Drive (SSD). In comparison to HDD, SSD provides faster transfer speed, smaller size, shock resistant, and efficient power consumption. The cost per volume ratio and a longer lifetime is the only advantages HDD currently holds. This transition with the storage devices has significant impact on digital forensics, especially with the preservation of evidence and data recovery analysis.

The literature reviewed in this chapter will provide a foundation and background for the research. Literature in relation to magnetic storage devices, NAND flash, digital forensics guidelines, and data recovery methods will be reviewed. Firstly HDD technology will be reviewed in Section 2.1. The historical background, components, mechanisms, and advanced features of HDD will be presented. Section 2.2 extends the foundation knowledge in storage devices to SSD. This review will clarify the architectural differences between the two, and links to Section 2.3 which will identify critical issues with SSD and available data recovery techniques. Section 2.4 examines a number of widely implemented digital forensic best practices and composes a possible best practice to mitigate the issue. Section 2.5 presents a summary of risks and issues identified in the literature review.

2.1 HARD DISK DRIVE

HDD is a magnetic storage device used where data is saved permanently. The name “hard disk drive” originates from a platter (disk with coating) being made of solid hard materials such as aluminium, glass, or ceramic. Historically floppy disks were dominant storage device in consumer market, but HDD has taken that position because of its larger capacity and faster access speed.

The primary HDD in a computer system is often called the “C drive” in Microsoft Windows operating system environment. Magnet field polarity is used on the platter and data is physically written on its surface by magnetic force. Data remains on the platter even if its power is switched off.

Due to HDD’s architectural design and harsh usage, HDDs has higher possibility of break down, in comparison with other computer components. Data are written on the spinning platter with extreme precision in nano-mechanical measurements, hence dust, vibrations or physical shock could easily interfere with the precision and damage could cause data to be no longer inaccessible.

In Section 2.1, details of HDD are reviewed to provide a foundation for computer data storage and architecture knowledge. Section 2.1.1 introduces HDD background and how it evolved over half a century. Section 2.1.2 describes HDD mechanism and how it operates. Section 2.1.3 explores HDD components in detail. In depth understanding of HDD components is essential to understand how deleted data can be recovered from the magnetic storage devices.

2.1.1 HDD Background

First HDD was developed by IBM in 1956. It was a part of IBM RAMAC system and fifty of 24" diameter platters were piled together, the whole size was almost equal to two larger size refrigerators together. The capacity of this HDD was 4.4 megabyte (Howe, 2008).

Increased popularity and demands for personal computer systems and especially processing speed and storage space became high in the 1990s. As Intel co-founder Gordon E. Moore published his prediction in 1965 paper known as "Moore's law", processing speed and storage capacity has been growing double each year (Hutcheson, 2009). Partly because the law has been adapted as a long-term guideline in semiconductor industry and set a milestones for research and development. The International Technology Roadmap for Semiconductor (ITRS) reports the trend continued for over half a century will diminish at the end of 2013, estimate speed and capacity will only double every three years (Sood, James, Tellis, & Zhu, 2012).

For the majority of users, including professionals working in IT industry, HDD is a storage device commonly used but used without the need for an understanding of its mechanism and architecture. Data recovery on HDD takes a great portion of computer forensics analysis time and it is hard to do so without knowing the device architecture.

It is proven that when a technology is understood well, data can be recovered in a way that even experts thought unachievable. For example, in a homicide crime case, a floppy disk was shredded into pieces neither law enforcement nor the manufacturer had a protocol to recover data residing on it. After several failed attempts, a forensic team taped all the pieces together on cardboard aligning the original tracks and spending just $131, and then recovered more than 80% of the data (Duffy, 2004).

Most HDDs share the same mechanisms, and the basic architecture has remained unchanged since the beginning of the 1980's (Kozierok, 2001). Therefore current HDD data recovery methods are fairly well established, and well proven in the digital forensics field.

2.1.2 HDD Mechanism

On HDD, data is recorded on a spinning round flat disk called *platter,* and an electromagnetic *head* reads and writes data stored on the platter surface. The head is attached to a tip of an arm with a suspension called a *slider*. The head, the arm, and the slider, three pieces together is called a *head assembly*, its position and movement is regulated by an *actuator*. The actuator controls the head assembly's arc of motion in order to access data in designated area on the platter surface. The head operates just above (often tens of nanometres) the surface without physically touching.

The platter surface is coated with magnetic material to store information by changing polarity with electromagnetism, and segmented into a concentric circular area called *tracks*. The tracks are radially divided into the smallest unit called *sectors*. Data is stored on sectors and traditionally each sector is capable of storing 512 byte for traditional HDDs and 4 kilobyte (symbol: KB, 4096 byte) for newer HDDs. When multiple platters are used in a HDD, platters are piled with the head assembly in between and the multiple tracks on each platter can be accessed simultaneously. The tracks that are physically located directly above each other are called a *cylinder* (Tech Juice, 2011).

The platters rotate from 4,200 rpm (revolutions per minute: number of times the spindle of a motor rotates in one minute) in energy efficient portable drives, to 15,000 rpm for high-end performance server drives (Blount, 2007). The spindle is connected to a dedicated motor called spindle motor.

Physical format and *Logical format* is required in order to save data on a HDD. The physical format is also known as *low-level formatting*, which sectionalises the platter surface into basic entities of tracks, sectors, and cylinders, by polarising the platter areas using the heads. The purpose of physical format is to prepare the platter surface to be written on and for manufacturer to identify *defective sectors*. Therefore consumers do not need to perform physical formatting as it is usually completed by manufacturer (Kioskea, 2013).

A Logical format is also known as high-level formatting, which creates a *file system* on the platters. The file system is required by an *operating system* to use HDD space to store (write) and to access (read) data, and it differs depending on operating systems. If multiple file systems are required on a hard drive disk, storage area can be divided with *partitions*. Each partitions are allowed to have own file system, this accommodates issues with multiple file systems and compatibility on single HDD (Kioskea, 2013). The main purpose of a file system is to define an allocation table (e.g. *File Allocation Table*, also known as *FAT* format type) to efficiently access data without a need of searching entire storage space. Logical format also contributes for efficacy by organising a group of sectors into a slightly larger unit called a *cluster*, which reduces the overhead of data structuring and allocation (Christensson, 2005). The file system does not allocate data to individual sectors but instead uses clusters.

2.1.3 HDD Components

A HDD consists of approximately 300 component parts. Some of key components are previously introduced. Further details of functionality, material used, and implication are discussed from Section 2.1.3.1 to 2.1.3.10.

2.1.3.1 Disk platters

All HDD contains one or more platters which in fact stores data on its surface, and a typical HDD has multiple platters for faster access speeds and increased storage volumes. HDDs for personal computer (PC) systems are made available in various sizes. These are categorized by physical size of the platter of 3.5 inch is commonly seen in desktops PC, where 2.5 inch or 1.8 inch are found in laptop PC. There were 5.25 inch HDDs used in 80's to early 90's but these are made obsolete due to smaller HDD form factors development.

Platters are composed of two main materials: a *substrate* material that forms the core of the platter providing rigidness, and *magnetic media coatings* which "holds the magnetic patterns that represent data" (Kozierok, 2001). Traditionally aluminium alloy was used as a substrate material but manufacturer's aspiration for higher densities achieved glass made material called a *glass ceramic composite*. Innovation of the glass ceramic composite platters, or commonly known as *glass platters*, halved the thickness of the conventional aluminium platters (Mueller et al., 1998). The weight of platter is a key factor for HDD life

span, lighter the weight is better. Lower heat conductivity brought another advantage to thermal stability. The aluminium platters suffered from expansion and contraction depending on temperature (Mueller et al., 1998). Most of the HDD manufactures are shifting towards glass platters as a standard substrate, especially in high performance models (Kozierok, 2001).

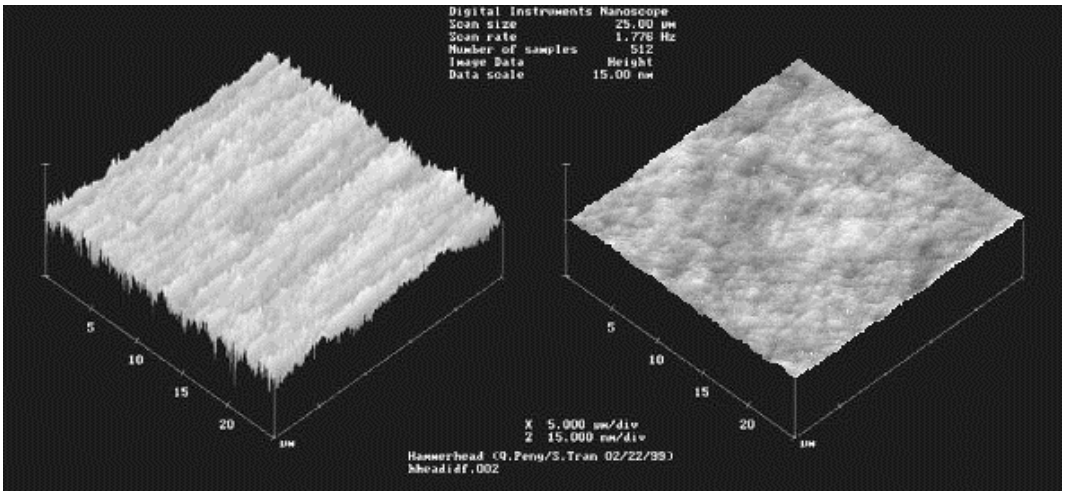

Figure 2.1: Hard disk platter viewed with a scanning electron microscope. The image on the left is surface of aluminum alloy, right is a glass. (Kozierok, 2001)

The platter surface needs to be exceptionally smooth and flat, because the platters spin and read/write heads float just above them. Older and slower HDDs had high fly heights of the heads and evenly flat surface was not a significant issue. Demand for higher densities and smaller form factor, the openings between the surfaces and the heads are closing (Kozierok, 2001). Image driving a car at a faster speed on a narrower road then the risk of collision or accident increases dramatically and the demand for a flat smooth surface will become essential. The head touching the surface is called a *head crash*, and this will stop HDDs to operate. Recovery in such event will be discussed in Section 2.3.

The magnetic material coating is a thin layer of magnetically retentive substance covering the platters surface, in which information is stored by altering magnetic patterns (Mueller et al., 1998). Traditionally *oxide media* was used to cover aluminium platters. “Oxide” means *iron oxide*, which is also commonly known as rust. Although rust is an effective compound for magnetic coating material, it is considered not an ideal marketing term and therefore manufactures rephrased it. For example, “high performance oxide media layer covering rigid core made of aluminium substrate”. This media is coated by spinning the platter at high speed, centrifugal force to spread the media from the centre to the outside of

the platter evenly. The coated surface is then polished and covered with another layer to protect and lubricate (Mueller et al., 1998). The finished oxide coated platter looks like a rusted disk. The oxide media coating is around 0.8 microns (μ, or 30 microinch) thickness (Kozierok, 2001). This type of media is also found in recording magnetic tapes, such as audio cassette, video, and Linear Tape-Open (LTO) data storage.

Oxide media has been used since 1955, the beginning of HDDs. Requirements from modern HDDs exceeds the limit of oxide media capabilities therefore it is replaced by thinner, harder, and more perfectly formed media called *thin-film media* (Mueller et al., 1998). As the name suggests, thin-film is thinner than its predecessor.by all means, oxide media was reasonably thin, but the new material is exceptionally thinner in comparison. It is also known as *plated media*, or *sputtered media* due to its requirements of distinctive application method to deposit such thin material onto the platters (Mueller, 1999).

Thin-film plated media is produced by employing electroplating mechanism which is used in vehicle chrome bumper or jewelleries (Mueller, 1999). The platter is dipped in a series of chemical baths, coating the platter with layers of metallic films (Mueller et al., 1998). The magnetic layer itself is about 0.078 μ (or 3 microinch) thickness (Mueller, 1999).

Thin-film sputtered media uses a vapour-deposition process method, applied from semiconductor manufacturing of silicon wafers, to "coat the platter with a layer of nickel phosphorus and then applying the cobalt-alloy magnetic material" in a continual vacuum-deposition process called *sputtering* (Mueller et al., 1998). This method achieves magnetic layer of between 0.026 to 0.052 μ (or 1 to 2 microinch) thickness, even thinner than thin-film plated media (Mueller et al., 1998). The sputtering method is then used again to mount another layer of extremely hard 0.026 μ thick protective carbon film. Sputtered media coating has the advantage of a more evenly flat surface than plated media.

Uniformly smooth surface allows the head to float closer to the platter surface, floating heights as low as 0.078 μ above the surface can be achieved (Mueller, 1999). The head and floating mechanism will be discussed in Chapter 3. Basically when the heads floats closer to the surface, "the density of the magnetic flux transitions can be increased to provide greater storage capacity" (Mueller, 1999). In addition, the greater intensity (or closer distance) of the magnetic field

during read/write process provides higher signal amplitudes, which leads to better signal-to-noise performance (Mueller et al., 1998).

Both plating and sputtering methods provide thinner, harder film of magnetic media on the surface. Harder surface protection increases chance of data survival in an event of head crash occurs while the platters are spinning at high speed. In fact, modern thin-film media is resistant to head crashes, where oxide media coating is much more likely to be damaged (Kozierok, 2001).

Although near perfect vacuum-deposition makes sputtering costly process, due to the increased demand for meticulousness in quality claims thin-film sputtered media as a primary method on current HDDs. The sputtering method results in the most desirable, thinnest, and resilient platter surface which can be manufactured commercially (Mueller, 1999).

2.1.3.2 Read/write heads

A HDD typically has one read/write head for each platter surface, and typical modern HDDs has four platters and eight corresponding read/write heads. These heads are *ganged* together on a single harmonized movement mechanism, but only one head can perform read or write at given one time (Kozierok, 2001).

The roles of read/write heads are vitally important for HDD operation as well as the overall PC systems. They are also the most expensive components of the HDD and they determine the performance of the read/write speed and HDD performance (Kozierok, 2001). However not many users have seen this sophisticated, critical component sealed inside the HDD chassis.

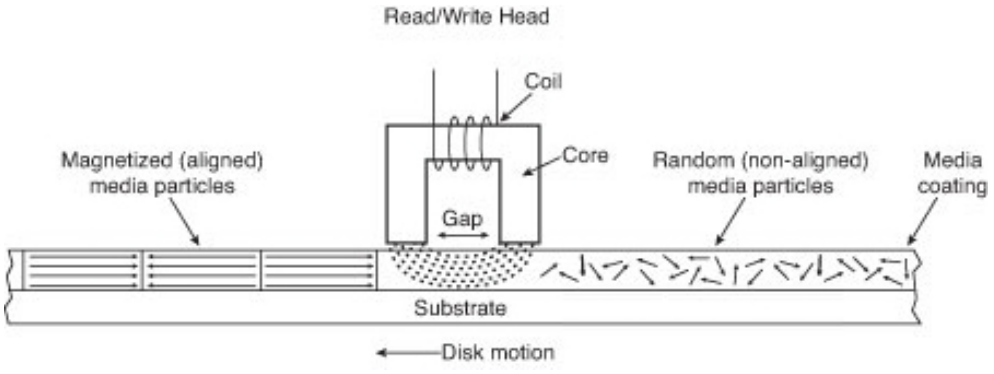

Figure 2.2: Illustration of how read/write head interact with magnetic information on the platter (Bestofmedia, 2011)

Mechanically, read/write heads are relatively simple. They are energy converters which they transform electrical signals to magnetic polarities (write), and magnetic polarities back to electrical signals (read) (Kozierok, 2001). The

read/write heads perform this conversion between electrical information and magnetic (Figure: 2.2), and bit of data is saved on the platter using an *encoding method* that "translates binary into patterns of magnetic flux reversal" (Kozierok, 2001). The primitive heads were simple electromagnet, an iron core with winding coils. Mainly four types of head schemes have progressed over the years.

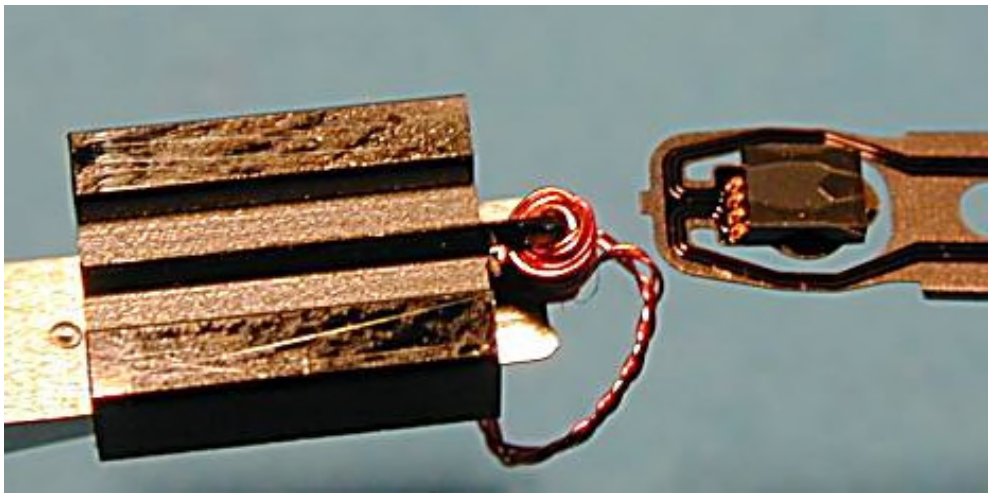

Figure 2.3: A graphic illustration of 15 years evolution of HDD head sliders. At left, a slider from a 40 MB 5.25" ferrite-head drive; at right, the slider from a 3.2GB, 3.5" MR-head drive (Mueller et al., 1998, p.740)

Ferrite heads are the traditional type of magnetic-head design, inherited from the original IBM Winchester drive and popular in low-end market during the 1980's (Bestofmedia, 2011). They are larger and heavier in comparison to newer heads such as thin-film heads, and therefore required higher floating heights to prevent a head crash (Mueller et al., 1998).

Metal-in-gap (MIG) heads are improved form of ferrite heads. The sputtering technique is used to meet demands for higher density recording by increasing resistance to magnetic saturation (Mueller et al., 1998). MIG heads replaced ferrite heads from the competition until early 1990's.

Thin-film (TF) heads are manufactured in much the same manner as the silicon wafers discussed in the platter magnetic material coating methods. A *photolithographic* manufacturing process can mass produce smaller size and higher quality product with cost effective way. Sophisticatedly smaller TF heads can float much lower height above the surface, produces a sharply defined magnetic pulse, and captures stronger signals (with less signal-to-noise interferences) from the surface due to its defined sensitivity and lower altitude

(Mueller et al., 1998). Once again, TF heads displaced MIG heads from the market and still remain strong in mainstream HDDs.

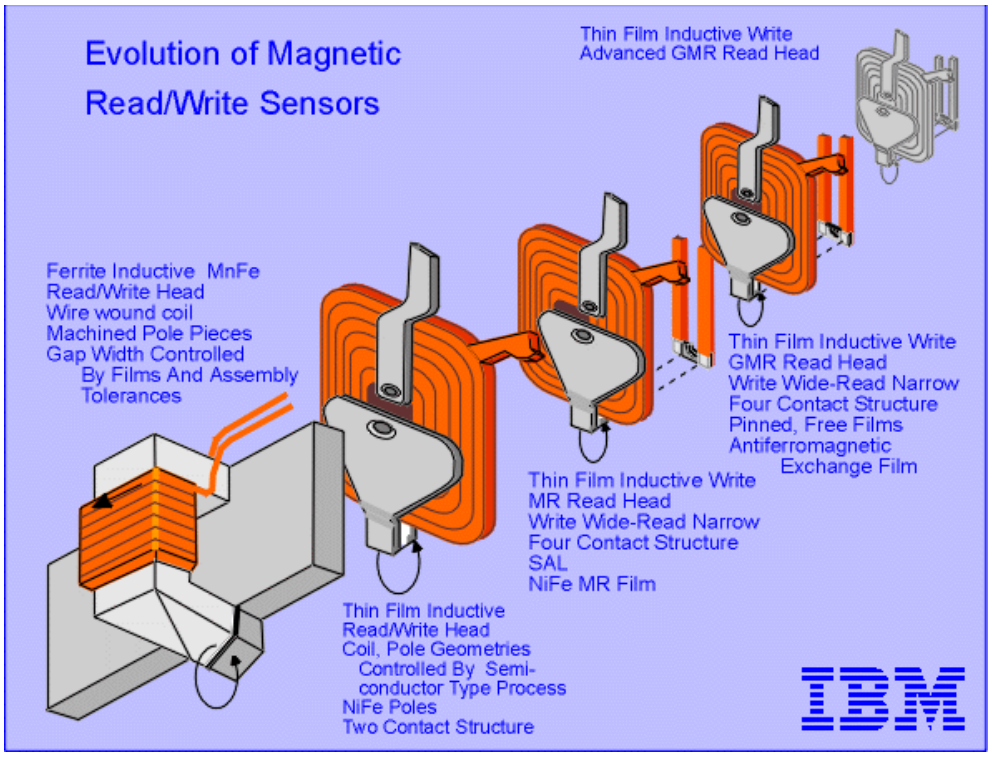

Figure 2.4: Summary chart showing the basic design characteristics of most of the read/write head designs used in PC HDDs. Original image © IBM Corporation (Kozierok, 2001)

Magneto-resistive (MR) heads are the most recent technology (Mueller et al., 1998). While conventional heads (such as ferrite, MIG, and TF, also known as *single-gap heads* because the same gap is used for both read and write) rely on principles of electromagnetic force to read and write information, MR heads use principle of *magneto-resistance* (Kozierok, 2001). Developed by IBM in 1991, a MR head act as a resister sensing current changes in resistance (Bestofmedia, 2011). The MR principle can only use for reading process and TF head must be accompanied for writing data. This apparently limited feature brought benefits of further optimization. Previously single-gap heads must share the same gap for both reading and writing and in other words performances are compromised to meet requirements for both. This is no longer an issue MF heads because two separate heads are assembled together, each heads can be independently tweaked for its task (Mueller et al., 1998).

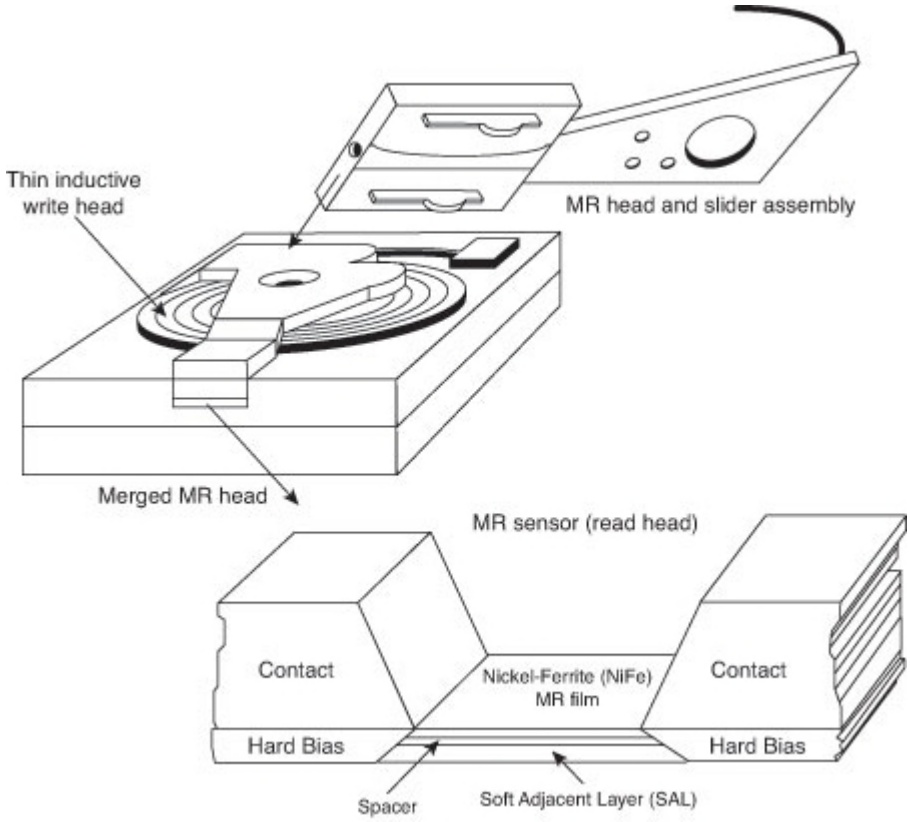

Figure 2.5: Illustration of MR head composition (Bestofmedia, 2011)

Negative sides of MR heads are that it is more costly to produce when compared with TF heads because of additional components and processes involved. But the design can reach at least three times more amplified reading ability than TF heads (Mueller et al., 1998). MR heads are also more vulnerable to stray magnetic fields due to refined sensibility, and requires better shielding (Mueller et al., 1998).

MR heads continue to evolve but the basic principle still stands without changes. IBM announced the first *giant magneto-resistive (GMR) heads* in December 1997 (Bestofmedia, 2011). In fact, GMR heads are smaller than MR heads but named GMR effect achieved even greater density. Furthermore, Hitachi has developed GRM heads using perpendicular current which supports areal densities of up to 1 Tbits/sq. This is called *current perpendicular-to the plane giant magneto-resistive heads (CPP-GMR)* and commercially debuted in 2011 (Bestofmedia, 2011).

Overall, the GMR head is the main player in the market. As the density became higher and contributed to greater storage volume on a platter surface, signals became weaker and exposed to additional vulnerability such as stray magnetic field interferences. In order to compensate such issues, an electrical

pulse amplification circuits were developed for improved electrical to digital signal conversion and an error detection and correction circuit must be integrated to "compensate the likelihood of errors" (Kozierok, 2001) due to the need for interact with feebler signals.

2.1.3.3 Head sliders

Modern HDD heads float over the spinning surface of the platter and perform read/write data without ever physically contacting the surface. This is one of distinguishing difference of HDD architecture from other conventional magnetic tape storage, such as audio cassette tapes and video cassette recorders (Kozierok, 2001). The tiny space between the heads and the platter is called *floating height* or *head gap* (Kozierok, 2001). The read/write heads are mounted on suspension *arms* that firmly pressure a body of material that supports the heads called *sliders*, to the platters while not spinning. This is to ensure position and floating height of the heads are properly maintained (Kozierok, 2001).

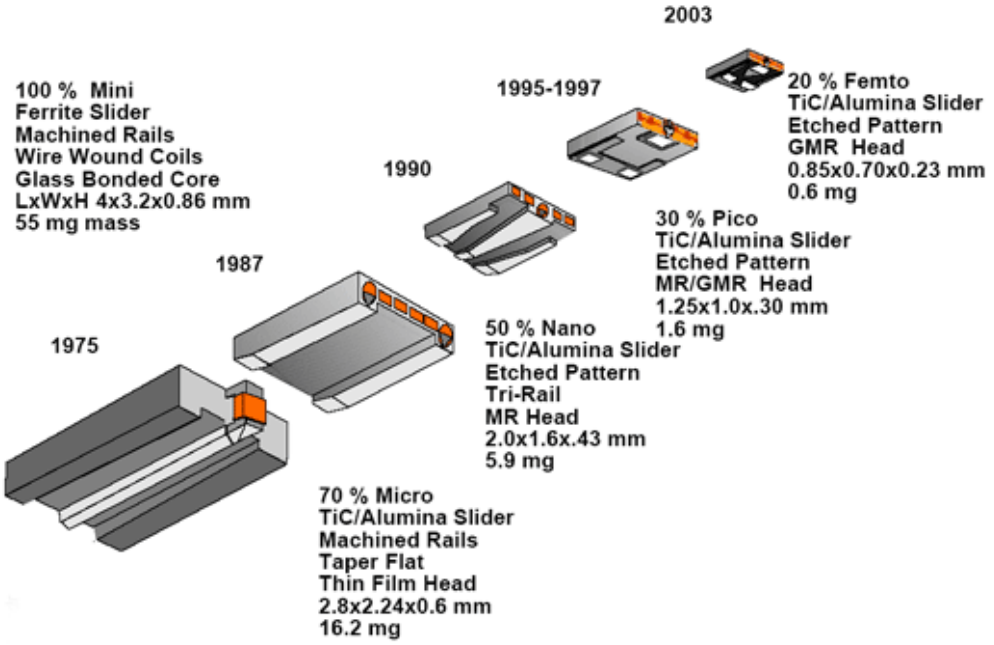

Figure 2.6: Evolution of HDD sliders (Brooker, 2005)

Imaging an airplane taking off or landing on a runway, landing gear with tyres is touching the ground not the airplane itself. The sliders act the same as the landing gear, carrying the head at the correct distance above the platter for reading and writing (Mueller et al., 1998). Once the platter begins to rotate for operational speed, the high speed creates airflow under the sliders and lifts them above the surface, just like airplanes on take-off from a runway (Kozierok, 2001).

The advance towards smaller and smaller form-factor requires for equally smaller slider size as well (Mueller et al., 1998). Traditional mini-Winchester slider dimension is about 4x3.2x0.86 millimetre (mm) in 1980's.

Table 2.1: The characteristics of the various types of sliders used in HDDs (Kozierok, 2001)

Slider	Year Introduced	Relative Size	Length (mm)	Width (mm)	Height (mm)	Mass Type (mg)
Mini	1980	100%	4	3.2	0.86	55
Micro	1986	70%	2.8	2.24	0.6	16.2
Nano (+ Pressure)	1991	62%	2.5	1.7	0.43	7.8
Nano (- Pressure)	1994	50%	2	1.6	0.43	5.9
Pico	1997	30%	1.25	1	0.3	1.6
Femto	2003	20%	0.7	0.7	0.23	0.6

As shown in table 2.1, the size of slider has been significantly miniaturised. Current the *femto-slider* is 20% of the original mini-sliders size and the weight (mass type) is only 1%. This reduces mass carried at the end of actuator arms and contributes for improved access seek time (Bestofmedia, 2011). Furthermore, smaller slider has even lesser contact area during slowing down and starting up the platter rotation, which reduces minor wear on the protective coated platter surfaces (Mueller et al., 1998).

Floating heights varied significantly in conventional sliders, depending on the velocity of the platter surface travelling below (Mueller et al., 1998). For example, outer platter surface has higher velocity and floating heights, and this phenomenon is undesirable in newer HDDs with the same bit density across the entire platter surface. When the same bit density or *zone recording* is achieved, the floating height is also required to be consistent with minimum variables for optimum throughput. The platter surface patterns are modified to adapt newer sliders, and special textured and manufacturing process enables the floating height to be stabilised, making desirable zoned recording HDDs (Mueller et al., 1998).

2.1.3.4 Head actuator mechanism

Recent trend shows manufacturers are reducing the number of platters, even from their flagship models. One reason for this trend is because having multiple head assemblies (assembly of head, slider, and arm) requires more complex tuning to meet high speed positioning with precision (Mueller et al., 1998). This is due to increased weight in each additional arm, and alignment of multiple heads. A

decreased number of platters simply mean a less complex mechanism, but also a trade-off of available storage space for lesser platter surfaces.

A device called a *head actuator* governs Head assembly's movement or positioning. This mechanism moves the heads in arc of motion, positioning them precisely above the target cylinder (the tracks that are physically located directly above each other) (Mueller, 1999). Modern head actuator uses *moving coil motor* or commonly known as *voice-coil*, almost the same mechanism used in audio speakers, to move the head arms and a closed-loop feedback (or guidance) system called a *servo system* to dynamically position the heads directly above desired cylinder (Goh, Li, Chen, Lee, & Huang, 2001). The voice coil actuator is not only thermally insensitive, but it performs at much faster speed with excellent reliability and accuracy (Mueller et al., 1998).

There is another noteworthy benefit of using a voice-coil actuator. While the platters and heads (or slider material to be accurate) are designed to endure brief physical contact in three occasions of starting up, spinning down, and power down, it is still better if any physical contact can be avoided especially where data is written. For this reason, most HDDs set a dedicated safe area called the *landing zone*, where no data is to be written there. The mechanism of guiding the heads to rest on this area is called *automatic head parking* (Mueller, 1999). The parking mechanisms initiate once a computer is shutdown or even in an event of sudden power loss. No special instruction or command is required for this process, hence automatic parking. In the simplest expression, the head arms have springs attached to both sides with one side weaker than another. On the stronger side of the spring, special space is made available called *park-and lock* position. As soon as a HDD loses power, the head arms are pulled towards the stronger spring side by overcoming the magnetic force of the positioner, and it is gently dragged to the designated parking space for its landing (Mueller et al., 1998). It is rare to find this ability in other head actuator architecture such as a *stepper motor*.

2.1.3.5 Air filter

Most of HDDs are sealed to prevent any particles, such as dust, going onto the platter surface. As shown in Figure 2.7, modern heads and platter surfaces operate in the world of microns, a typical dust particle on the surface will collide head to head (read/write heads) instantly and cause control problems. HDDs are sealed

but not airtight (Mueller, 1999). It requires specific tools and typically the warranty will be void if the seals are broken. The reason for not being airtight is to accommodate air pressure changes known as *pressure equalisation*. This is because the air pressure within the sealed HDD chamber is used to sustain optimised floating height of the heads, the pressure equalisation is required to stabilise air pressure (Mueller, 1999).

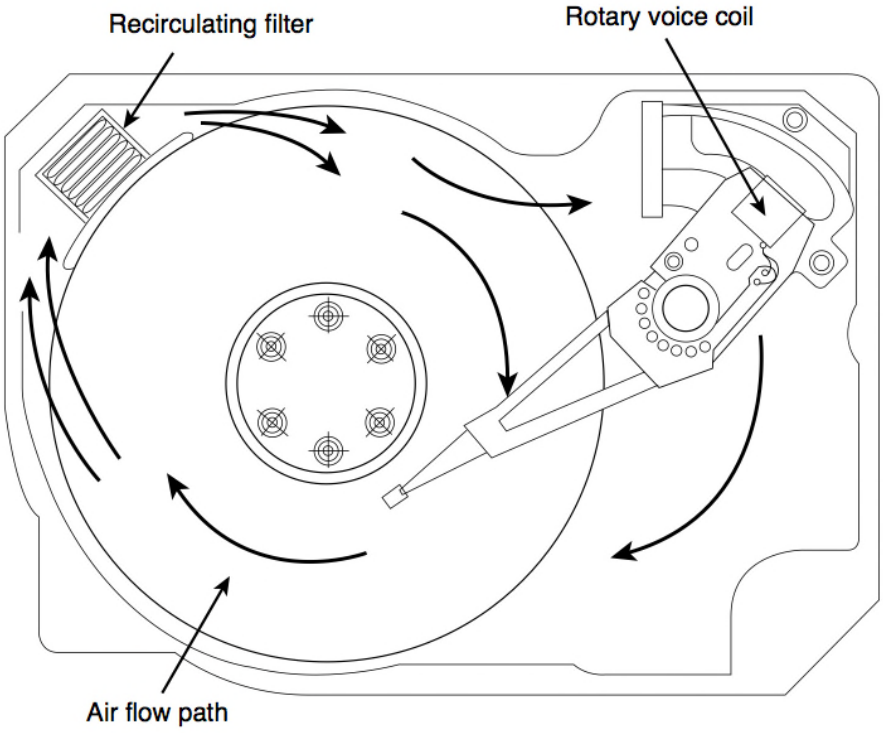

Figure 2.7: Location of HDD air filter and airflow (Mueller, 1999, p.607)

There are two permanently sealed air filters built inside typical HDDs, but these drives do not circulate air exhaustively (in and out). One is called *recirculating filter*, responsible for capturing any particles produced during standard HDD usage (Figure 2.7). A HDD is manufactured in cleanroom and it is tightly sealed semi-permanently. However the platters are rotating at high speed, head assembly is floating only 1 µ above the surface, constantly swinging in arc motion searching for desired cylinders to read or write electromagnet polarity. What would be the odds for chipping particles smaller than a micron? It makes good sense to have a filter to capture such particles for increased reliability and longer product lifetime. Another filer is called *barometric* or *breather filter* that is located on a *vent hole* inside the chassis (Mueller et al., 1998). The vent hole is where pressure is adjusted and is easily found with a warning sign stating not to cover the hole.

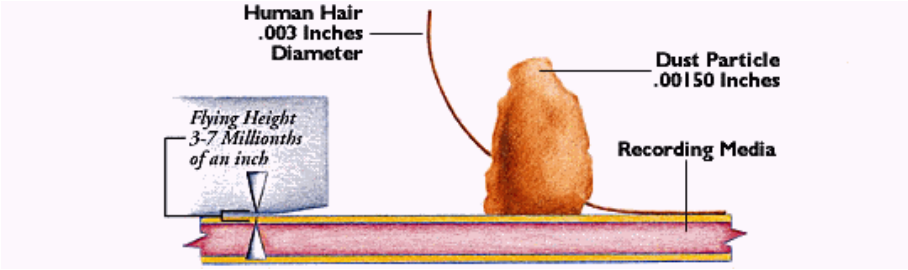

Figure 2.8: Comparison of average header flying height against a typical dust particle (Kozierok, 2001)

Typically manufacturers apply the barometric filter to prevent all particles larger than 0.3 µ while air bleeds through the vent. This meets the specifications for the cleanliness inside the HDD chamber in general.

The drawback of pressure equalization is that even finer barometric filters unable to prevent moisture entry. In any electrical components, humidity causes serious issues, such as condensation (Mueller et al., 1998). Most manufacturers have specific acclimation process for a HDD in new environment (Mueller et al., 1998).

Table 2.2: HDD Environmental Acclimation Table (Mueller, 1999)

Table 12.8 Hard Disk Drive Environmental Acclimation Table

Previous Climate Temperature	Acclimation Time
+40°F (+4°C)	13 hours
+30°F (–1°C)	15 hours
+20°F (–7°C)	16 hours
+10°F (–12°C)	17 hours
0°F (–18°C)	18 hours
–10°F (–23°C)	20 hours
–20°F (–29°C)	22 hours
–30°F (–34°C) or less	27 hours

Table 2.2 suggested acclimation time and temperatures. In a plain context, if a computer or portable HDD was taken out from cold area, such as boot of a vehicle in winter to inside a room, then the drive should not be used for at least 13 hours for acclimation.

All commercially available HDDs are equipped with these two permanently sealed air filters (Mueller, 1999). Special airtight sealed pressurised HDDs are made available for operating in high altitude (3000m or above will cause insufficient floating height due to lower air pressure) (Blattau & Hillman, 2004).

2.1.3.6 Spindle motor

The platters are rotated with a motor known as *spindle motor*, which is directly attached to the platter without any belts or gears. This mechanism is called a *direct drive method*, and provides essential characters such as noiseless, tranquillity, and stillness (Mueller, 1999). Imagine taking notes in a quiet room compared with doing so on public transport. Common speeds of rotation are 4,200 rpm (revolutions per minute), 5,400 rpm, 7,200rpm, 10,000rpm and 15,000 rpm. The faster rpm stands for better HDD performance and is usually more expensive. It is also important to maintain consistent speed of the rotation and an automatic control circuit precisely manages this (Mueller, 1999).

In relation to the spindle motors, few types of bearings are used between the spindle and motor coil. The bearings are used in spindles of HDDs, and orthodox spindles used *ball bearings (BB)*. BB suffered from wears due to continuous friction, and generated high frequency noise and vibration (Wilcoxon, 1994). Developments are made to improve but unable to resolve the issues.

Figure 2.9: Cross section diagram comparison of different bearing types (Aerocool, 2012)

Modern bearings use *fluid dynamic bearing (FDB)*, basically balls are removed and mucilaginous fluids are used to fill the gap as lubrication (Figure 2.9). Once the spindle reaches optimum rpm, it produces a longer life, low noise, anti-shock (no vibration), and a higher precision performance. However FDB consumes more power than the BB, and may not be able to reach sufficient rpm speed under extremely cold conditions due to the characteristics of the mucilaginous fluids (similar to vehicle engine oils in arctic regions) (Hashimoto, Ochiai, & Sunami, 2012).

2.1.3.7 Logic board

In the early days, primitive HDDs had a separate logic controller to perform all actions. Therefore the logic controller was generalized to support wider range of products. As technology progressed rapidly, each component became more complex and demand from customers drove further improvement to their performances (Kozierok, 2001). In the mid-1980's, due to the high demand for custom made logic controllers and miniaturisation of semiconductors, HDDs have integrated logic circuit board as a *logic board* (Mueller, 1999). The integration of logic board's ability to specifically tune each component produced faster speeds but the use of a separate logic controller became impractical (Kozierok, 2001).

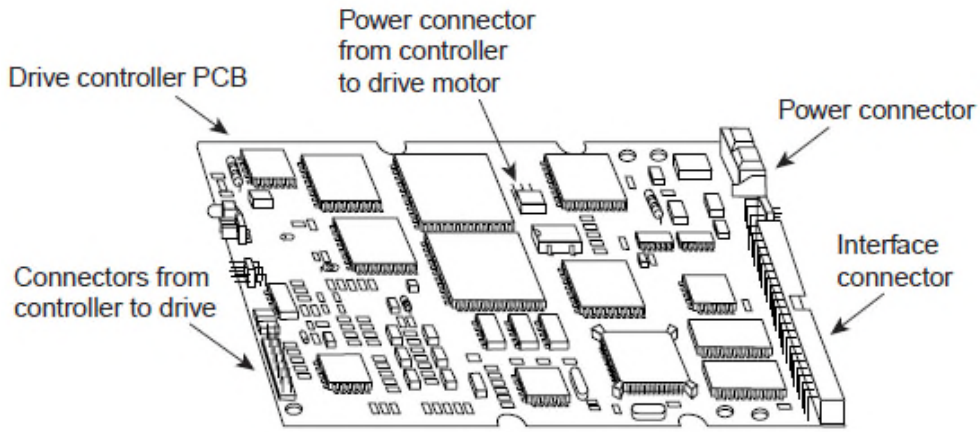

Figure 2.10: HDD logic board (Mueller, 1999, p.611)

Modern logic boards contain a microprocessor and internal memory. It is almost like a smaller PC embedded specifically for a HDD. This control circuitry manages the accuracy of the spindle motor, the actuator movements, read/write head process including electrical-to-magnetic signal conversion, signal amplification, power regulatory, and additional features for improved performances like internal cache optimization (e.g. pre-fetch) (Kozierok, 2001).

Interestingly, issues with logic boards are common cause of HDD failure, not in the mechanical components. The logic boards are typically mounted on the outside of the HDDs with screws and often openly exposed. It can be replaced easily at home, and replacement boards are widely available from manufactures. If the same HDD is available, a faulty drive can be tested by simply swapping the logic boards (Mueller, 1999).

2.1.3.8 Cables and connectors

There are two main types of connectors used on a typical HDD called a *power connector* and an *interface connector*. While integration of logic boards are standardised, there is still an external controller the HDD requires (Kozierok, 2001). The difference between an old logic controller and a new external controller is that an old controller manages all the components within the HDD and new controllers provide an *interface*, only acting as a communication link between the HDD and the system (Mueller et al., 1998). In 1990's, Small Computer System Interface (SCSI) and Integrated Drive Electronics (IDE) for general consumers were the popular interface preferences. IDE is then standardised under *AT Attachment* (ATA).

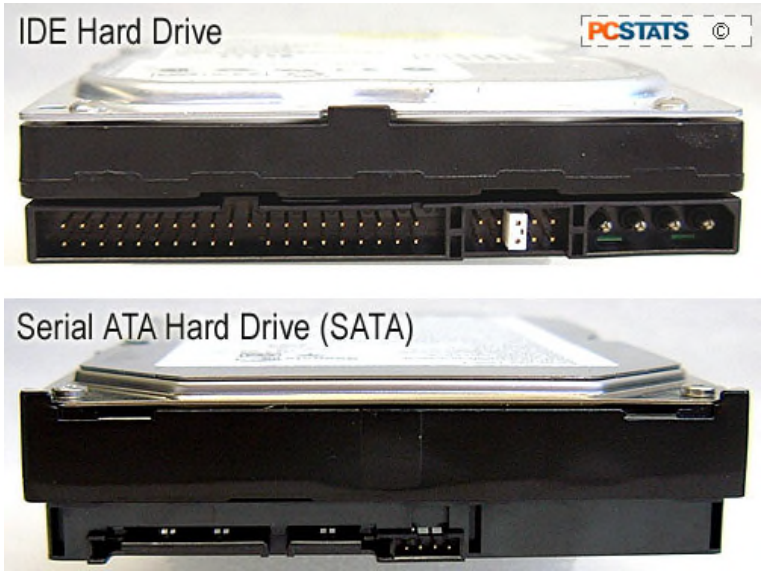

Figure 2.11: Pictures of IDE and SATA drives with different interface connectors (PCstats, 2006)

Serial-ATA (SATA), successor of ATA, is the most commonly used interface type in consumer or desktop markets at the present time, where *Serial Attached SCSI (SAS)* or *Fibre Channel (FC)* are preferred as an enterprise solution for their superior performance and expandability. HDDs are accessed through one of the interface types. A host bus interface adaptor commonly known as *south bridge*, is typically integrated on PC mainboards and a data/control cable is used to connect between them (Karbo, 2011). Each HDD also has an additional power interface which usually directly connected to a main power supply unit. With a modern HDD, the shape of power interface is unique, therefore an incorrect cable being plugged in can be avoided.

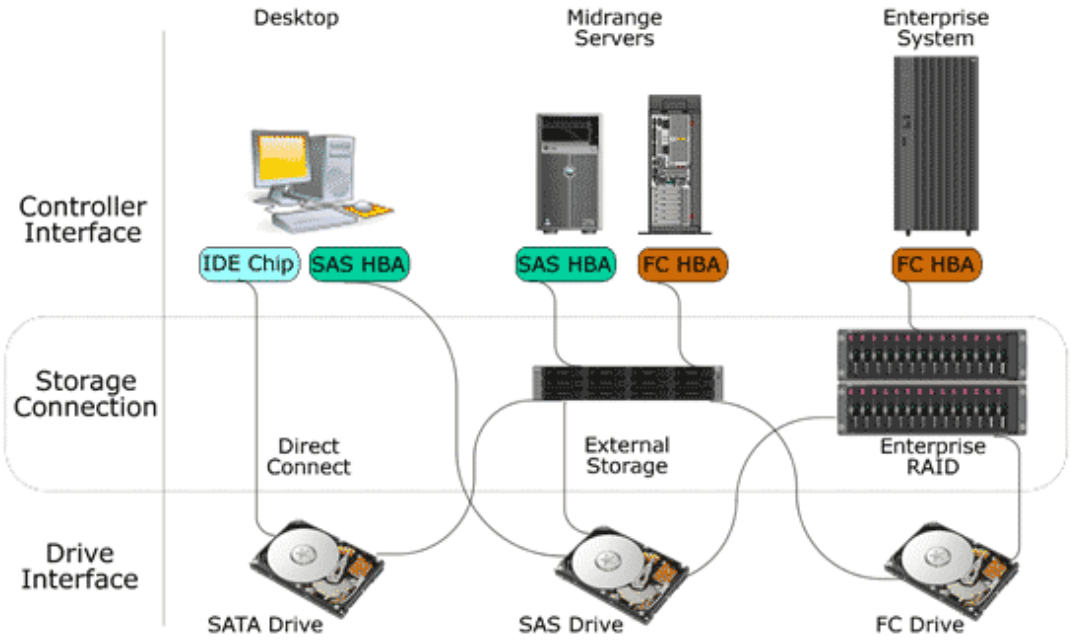

Figure 2.12: Disk and Storage Networking Deployment Scenarios (Fellows, 2007)

2.1.3.9 Mounting chassis

The internal density of HDDs continue to intensify, all the components became considerably sensitive due to its tauter assembly and requirements for micron-level of precision. This phenomenon also influenced the packaging of HDDs, and became crucial against external factors such as contamination from outside atmosphere, shock, noise, in order to maintain its reliability (Mueller et al., 1998). The packaging is consists of two main parts called a *base casting* and a *cover*.

As mentioned previously, most HDDs are not air tight, but sealed with a vent hole for pressure control purpose. Some manufacturer covers the base casting with a coating to eliminate potential risk of particles being sealed within the HDD chamber (Kozierok, 2001). Every component except the logic board goes inside the chamber. The sealed chamber should never be opened, not only it will void manufacturer warranty but particles can't be seen with human eyes could easily sneak into the chamber, defeating the whole purpose of the chamber being explicitly sealed. The opened chamber will not instantly cause HDD failure, but the risk of immediate malfunction is almost unavoidable. If for any reasons the chamber must be broken and opened, then a dust-free room or container called *clean room* must be used to prevent contamination. For the best computer forensic practice, a use of Class 100 (a clean room with no more than 100 particles larger than 0.5 µ per cubic foot area) will survive a legal attack (Mueller, 1999).

2.1.3.10 Configuration items

Conventional HDDs had configuration metal clips called *jumper*. It was predominantly used to determine which HDD is used by the system. There are several variations depending on interfaces and drives, but modern HDDs, such as SATA drives, do not require such manual configuration anymore (Mueller, 1999).

2.2 SOLID STATE DRIVE

While HDD steadily expands its capacity and performance, current trend such as cloud computing, mobile computing, and high-end users demand faster speed for performance and less energy consumption for efficiency. *Solid State Drive* (SSD) has history since 1950's and current NAND flash-based SSD emerged since 1995 developed by M-Systems.

NAND flash successfully retains the data in the memory cell without batteries unlike other *Random Access Memory* (RAM). Unlike HDD, as its name represents, SSD does not have any mechanical (dynamic) components and the structure is sophisticatedly simple. It didn't take too long for the SSD to become the alternative to HDD, delivering faster random access and transfer speed, less power consumption, improving cost performance, and fits in smaller form factors such as USB thumb drive (Micheloni, Marelli, & Eshghi, 2013).

A term NAND stands for "Not And", a binary operation in electric logic gate, and NAND Flash was invented by Toshiba in 1989 (Tal, 2002). Note that both USB drives and SSDs use the same Non-Volatile NAND Memory (NVM or NAND), however the quality of NAND used, as well as the controller and interface integrated make a distinctive difference between a simple USB drive and SSD used in enterprise blade servers (Seagate, 2011). In Section 2.2, the anatomy of SSD is discussed to distinguish the peculiar differences between HDD and SSD.

2.2.1 SSD Components

A SSD is a self-contained system consists of four main components (plus 2.5 inch enclosure) and all components are soldered on a *printed circuit board* (PCB). An image of SSD is shown in Figure 2.13. NANDs are organised in a different channel line layout. Technically each NAND has equivalent speed and an increase in numbers of parallel (concurrent) accessing will increase the speed. Currently

the majority of SSD controllers support 16 channel parallel access, and in theory enabling 32 channels should double the speed (Kitagawa, 2011).

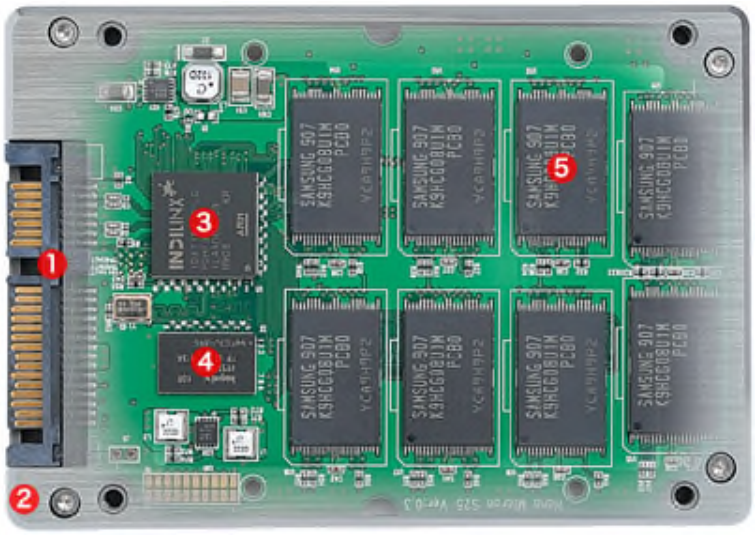

Figure 2.13: Anatomy of SSD components layout - 1.Interface 2.Enclosure 3.Controller 4.Cache 5.NAND Flash Memory (Kitagawa, 2011)

In addition to NAND and a controller, external DC-DC converter can supply internal power, filter capacitors can stabilise the power supply, and quartz can provide more precise clock (Micheloni, Marelli, & Eshghi, 2013). A fast *Double Data Rate (DDR) RAM* is generally used for cache. During write access, data is stored in the cache then transferred to the NAND. This is to prevent NAND wearing and also speed up the performance (Micheloni, Marelli, & Eshghi, 2013). NAND wearing and other issues will be discussed in later part of this chapter.

2.2.2 Non-Volatile Flash Memory

Semiconductor memories can be categorised into two types, previously mentioned RAM and *Read Only Memory (ROM)*. Distinctive differences are that RAM is unable to retain data without a power supply but allows random read/write access, while ROM holds data forever but is unable to write additional data. NAND belongs to the third category "NVM", which can retain data without power supply and the contents can be electrically altered (Micheloni, Marelli, & Eshghi, 2013).

NAND Flash Memory became familiar to consumers in 1990s, when portable products such as digital cameras, MP3 players, and USB flash drives grew to be popular. Although *NOR flash memory* (developed with NAND at the same time) is not covered in this research, it is worth mentioning that NAND and

NOR flash memories can be distinguished by looking at "how the memory cells are organized in the memory array" (Micheloni, Marelli, & Eshghi, 2013).

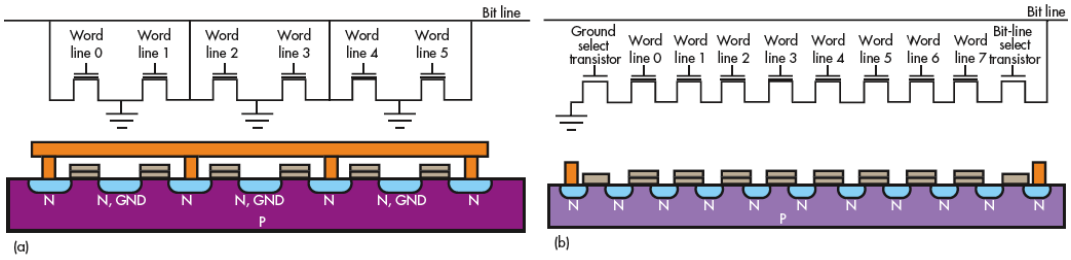

Figure 2.14: NOR (a) and NAND (b) flash memory, showing variations of bits per area between the two (Harris, 2007)

Due to the smaller size and the lower manufacturing cost, NAND flash which connects certain number of transistors, known as *floating gate (FG)*, in series are used (Kleinert & Leitner, 2008). FG transistor is consists of two overlapping gates, floating gate which is surrounded by layer of oxide to trap electrons, and *control gate* sits above the oxide layer encapsulating the FG and oxide.

This isolation gate mechanism is guaranteed years of charge retention (Micheloni, Marelli, & Eshghi, 2013). Injection and drainage of electrons are called *program* and *erase (P/E)*, and these operations modify the voltage threshold ($V_{TH)}$) of the memory cell. When FG's voltage is higher than V_{TH}, the cell is "1", otherwise it is considered "0".

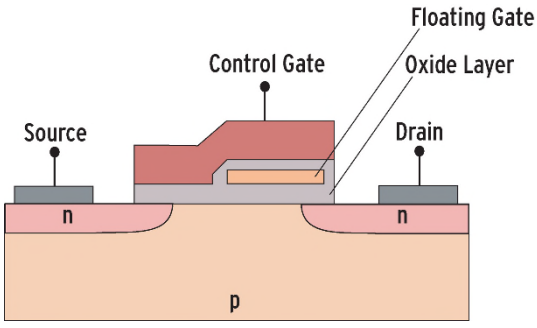

Figure 2.15: The oxide layer surrounding the floating gate prevents the electrons from escaping (Kleinert & Leitner, 2008).

2.2.3 NAND Array

Voltage of floating-gate transistors are used as memory cells for SSD. In *Single Level Cell (SLC)* devices, each memory cell stock one bit of data; where *Multi Level Cell (MLC)* devices stock 2 bits per cell (Micheloni, Marelli, & Eshghi, 2013). Some manufactures are already producing 3 bits per cell (*8LC* or *Tri Level Cell*) and 4 bits per cell (*16LC* or *Quad Level Cell*) devices. This fundamental

difference was essential to lower the NAND manufacturing cost; however there are few trades offs to the consumers. In comparison with SLC, multiple-level-cell (e.g. MLC, TLC, QLC) devices suffer from lower endurance, life expectancy, and performance. While all NAND flash memories provide unlimited read performance, they have limited *write (P/E) endurance*. When a memory cell reaches the limit, the failing rate will increase exponentially and the device will retire the cell from duty. The number of P/E cycle is determined by the silicon size and the bit density, "smaller lithography size and higher bit per cell will lower the write cycle" (Nguyen, 2013). Higher bit density suffers from increased error rates and retries attempts, and this leads to lower performance. In practice, SLC are used in critical infrastructures where higher performance and reliability is required, and MLC are commonly used in consumer products.

A *logical page* is the smallest unit made available for read/write process in NAND flash memories, and a *block* which consist of multiple pages, is the smallest unit used for erasing (Figure 2.16). For example current high capacity device consists of multiple *Logical Unit Numbers (LUNs)*. Each 32GB SLC LUN has 4096 blocks, each block containing 128 pages, and each page has up to 8640 bytes (Cooke, 2011).

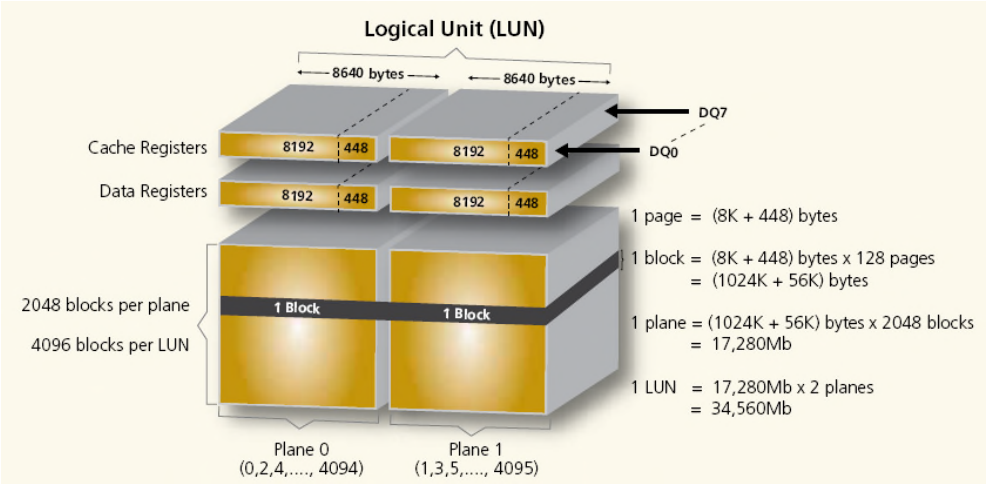

Figure 2.16: Sample 32GB SLC NAND memory array logic organization (Cooke, 2011)

A logical page is made up of *strings* of memory cells which divided into *main area* (data) and *spare area*. Each cell is connected through *wordlines (WLs)* to form a string, and pages consist of cells belonging to the same WL (Micheloni, Marelli, & Eshghi, 2013). The main area can be configured to 2KB. The spare area is an extra space for additional functionality such as *Error Correction Code*,

and in this example 64 bytes of spare space is made for every 2KB of main area. Details of additional functions are discussed separately in later part (Cooke, 2011).

2.2.4 NAND controller

A *NAND flash memory controller* is a core responsible for SSD performance and reliability. Two fundamental tasks are assigned to the memory controller:

1. Acting as a logistic manager, providing the best suitable protocol and interface to both host and NANDs;
2. Acting as an operation manager, maximising performance by efficiently handling data, maintaining data integrity and evenly distribute workloads (Micheloni, Marelli, & Eshghi, 2013).

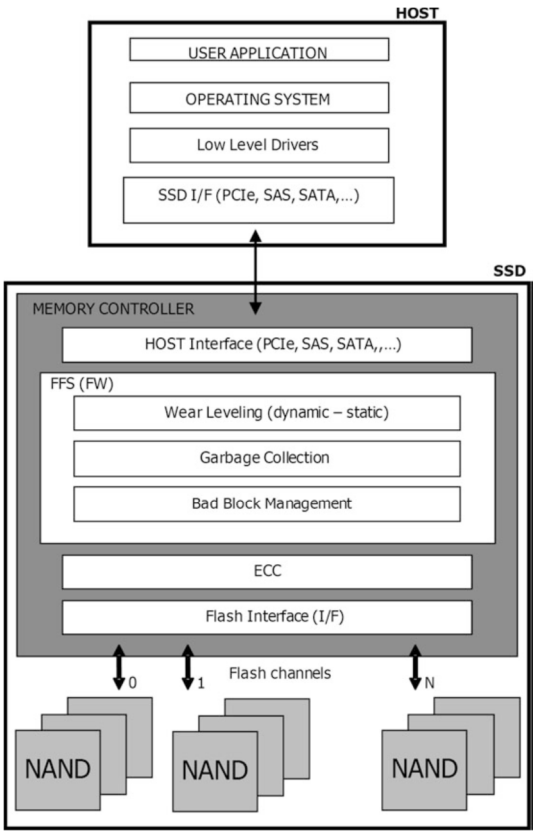

Figure 2.17: High level architecture view of NAND flash memory controller (Micheloni & Eshghi, 2013, p.27)

A high level diagram of common memory controller is shown in Figure 2.18. Due to the fact that performance of a memory controller strongly governs SSD performance, manufactures are striving to develop better products by adding unique functionalities, but the basic design remains similar and consists of three parts.

Proceeding from the host to the NANDs, the first part is the *host interface*. The majority of traditional HDD interface protocols are IDE (e.g. ATA133), SATA, SAS, or PCIe and SSD is design to be compatible with the same industry standards, although it is rare to find SSD with IDE interface because transfer speed SSD is capable of easily exceeding the IDE and therefore has a bottle neck for its true performance (Micheloni, Marelli, & Eshghi, 2013).

The second part is the *firmware* or also known as *Flash File System*, which enables the SSD to be used as HDD and equipped with sub-layers of functions (Kawaguchi, Nishioka, & Motoda, 1995). These functions are the main components and discussed separately. Other than these functions, the firmware manages the file system transition. SSD stores lists of sub-sectors which constitute a file, and file allocation table requires these lists to operate functions such as read, write, modify, and delete (Micheloni, Marelli, & Eshghi, 2013).

The third part is *Error Correction Code (ECC)* or also known as *Error Detection Code*, which corrects *bit errors*, thus has direct correlation with P/E endurance. Depending on the SSD, bit errors are logged and corrected, then liaised with the *Bad Block Management* one of the firmware's sub-layer functions. *Reed-Solomon coding* and *BCH coding (The acronym BCH comprises the initials of these inventors' names)* are commonly use ECC algorism. Note that the ECC is often treat as a part for the firmware functions, but generally ECC is a separate hardware and a part of the memory controller (Herth, 2011).

Returning to the discussion on the firmware's core functions, as shown in Figure 2.17, *Wear Levelling, Bad Block Management,* and *Garbage Collection* are unique to the SSD and are never required on traditional HDD devices (Kitagawa, 2011).

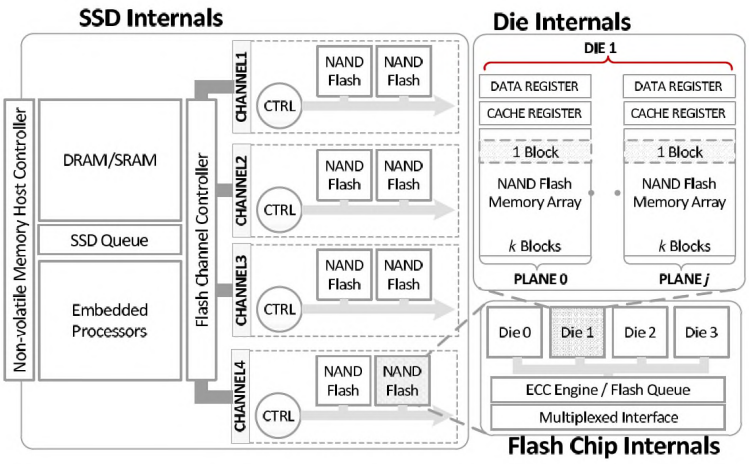

Figure 2.18: Modern SSD internal architecture. Note that an I/O request can be simultaneously served by many internal resources, which is one of the important characteristics of SSDs (Jung & Kandemir, 2013)

2.2.4.1 Wear levelling

It is eminent that there are files that change regularly and some files remain as is for longer time than others (e.g. system files). Regardless of bit density, ultimately memory cells have a limited P/E cycle and eventually become obsolete and removed from the recycle pool. The wear levelling was introduced to evenly maintain the cell usage to avoid heavy P/E cycle usage on certain cells. This is achieved by a simple monitoring of the aging of each page and block as little and as consistent as possible (Micheloni, Marelli, & Eshghi, 2013). *Dynamic Wear Levelling* is effective for usual operation while *Static Wear Levelling* is required for files with least changes are made to it.

2.2.4.2 Bad Block Management

The memory cells have life expectancy and having obsolete cells are inevitable. When a block has an obsolete cell, it becomes a *Bad Block* and added to a map for maintenance. The map is created during factory initialisation of the SSD which means every NAND flash memory has certain number of bad blocks at the point of manufacturing. The map is updated as soon as bad block is detected in order to retain data reliability (Micheloni, Marelli, & Eshghi, 2013).

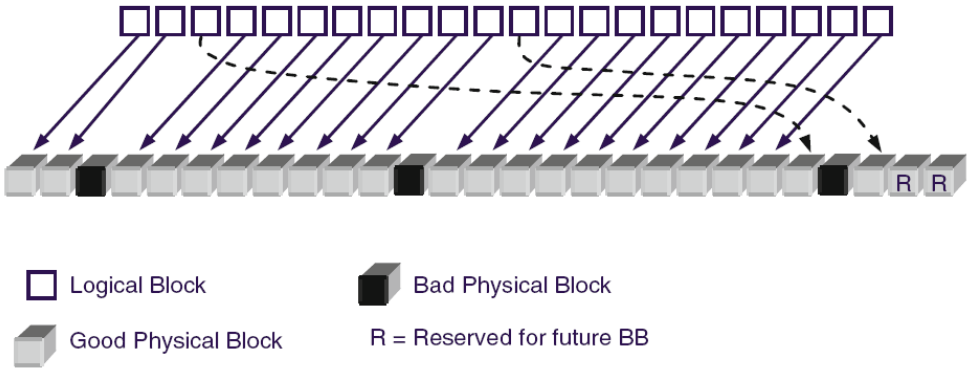

Figure 2.19: Bad Block Management (Micheloni, Marelli, & Eshghi, 2013)

2.2.4.3 Garbage Collection

There is a misconception that SSD does not require *defragment*, where HDD suffered from fragmented sectors (i.e. imaging a phone number written everywhere on a paper without sequence, takes longer time to read) and frequent defragmentation was suggested to restore its performance (Kitagawa, 2011). This is half correct, and the truth is that SSD never necessitate defragmentation in order to improve the read performance, but when the volume of free space drops below the set threshold, which impacts the wear levelling function, SSD initiates the garbage collection in the background. The process similar to defragmentation which copies the latest data, eliminates invalid sectors including the deleted files, then frees up the original block (Micheloni, Marelli, & Eshghi, 2013). This function is completely automated and transparent to the users for efficiency, which also means there is no way of manually stopping this process for salvaging deleted files.

2.3 DATA RECOVERY METHODS

Physical components, material, basic architectures for HDD and SSD are discussed in Section 2.1 and Section 2.2. In order to discuss data recovery methods, Section 2.3.1 elaborate on the read/write process on a magnetic storage device, and then is followed by data recovery methods in Section 2.3.2 and 2.3.3.

2.3.1 Read/Write Process

In Section 2.2, physical mechanisms of how data is read/write with HDDs are discussed. Basically electromagnetic polarities of the magnetic coating on the spinning platters are changed by the head assemblies which float above the surface and controlled by the actuator. A read operation means the polarity (+ or -)

of desired sectors are read and converted into a binary (0 or 1). A write operation means to locate an available sector on the surface and converting a source binary and change the polarity of the sector accordingly.

When data is retrieved from a HDD, a command is issued to the operating system. It then determines corresponding head number, cylinder, and sector information where the data is physically stored on the HDD from a source called *file allocation table* (Coughlin, 2008). The operating system transfers this information to the logic board which dictates the HDD operation. The spindle start to rotate if not spinning already, the actuator swings the head assembly to position itself over the correct track and the head waits for the target sector (or cluster) to circulate right beneath the scanner (Coughlin, 2008).

When the desired cluster arrives beneath the scanner, the contents (data) of the cluster are pre-amplified for reading and converted into binary. The binary data is temporary stored in a space called a *cache*, which is located on the logic board, dramatically reduces loads from the operating system and improves access speed performance (Mueller, 1999). Contents of the cache are registered in a *buffer*, and provide further efficiencies for data retrieval between the system and the HDD. Finally the interface controller releases the binary data from the cache to the *random access memory (RAM)* on the computer system mainboard to be used by the operating system or an application (Coughlin, 2008).

Writing data on a HDD is similar to data retrieval but in reversal order. A block or frame of data flows into a cache and special mathematical error detection technique called *cyclic redundancy checking (CRC)* is performed for error correction. The operating system is responsible for determining a storage cluster; if the data is new then available clusters must be assigned. Depending on the data size, but typically if the data is larger than single cluster capacity, then the operating system instruct the HDD controller which cluster to begin writing the data with. If consecutive clusters are not available, then the heads swings between assigned clusters until whole data is stored (Coughlin, 2008).

The data cannot be written until the desired cluster rotates and arrives beneath the header. This waiting time is called a *latency* (IBM, 2001). Faster the platter rotates then shorter the latency will become. Thus the HDD performance can be significantly measured with its rotation speed (rpm). When the cluster reaches beneath the header, “a pattern of electrical pulses representing the data

passes through a coil in the writing element of the recording head, producing a related pattern of magnetic fields at a gap in the head nearest the disk" (IBM, 2001). When the magnetic orientation of the desired cluster area on the platter is altered it now represents the data (IBM, 2001).

The popular Microsoft Windows operating systems use fixed cluster size, and often clusters are not filled completely and leave some unused space called a *slack space* (Medlin & Crazier, 2010). For example 2gigabytes (GB) FAT16 partition uses 32kilobytes (KB), while the same capacity HDD with FAT32 partition uses 4KB. Benefits of selecting different file system is not relevant and will not be covered in this research, but up to 15% of space can be wasted due to the slack space phenomenon (Mueller, 1999).

2.3.2 Data Erasion and Recovery

The logical steps of data recovery and data deletion require specification. As previously mentioned, a typical Windows operating system writes data in a cluster size of 512 bytes on the platter surface. Due to the Windows file system format design, the cluster size is fixed and therefore often 512 bytes cluster space is not completely filled and leaving some spare space known as *slack space*. Physical location of each data's whereabouts are logically stored the path on a *file allocation table*. Once files are placed in the recycle bin, files are still recoverable by any user because the files are logically hidden from the user interface and not physically removed from the platter. When files are deleted and even removed from the recycle bin, the operating system erases the path from the file allocation table and label the clusters as free space while physically the data is entirely intact in the clusters. In addition, a term formatting HDDs usually means the file allocation table is rebuilt and again all used clusters remain physically intact (Medlin & Crazier, 2010).

According to NIST 800-88 Guideline for Media Sanitization (Gutienez & Jeffrey, 2006). there are various methods to securely destroy data from media known as erasing or wiping. Depending on the storage volume and performance, this procedure can be time consuming. When residual magnetism is concerned then the United States Department of Defence (DoD) 5220.22-M data sanitization method passes wiping process up to 7 times (Sawyer, 2006).

The reason for why residual magnetism shouold be concerned is discussed later in this section. Advantage of this method is that HDD remains functional and all spaces are overwritten with zero or random numbers. If further assurance is required then the guideline suggest *degaussing* and *destruction* (Gutienez & Jeffrey, 2006). Degaussing is a device that generates strong magnetic field and efficiently purge data from magnetic storage devices. Destruction method involves disintegration, pulverization, melting and incineration. Both methods will disable the device permanently, the media should withstand any attempt of data recovery (Gutienez & Jeffrey, 2006).

Various tools are made available to recover data from functional HDDs. Typically it is good practice to make two verified forensic copies (master and backup) of the original HDD of interest, and conduct investigation only on master forensic image. Commercial software, such as Guidance EnCase, can analyse not only available files but also recover and search for deleted items (even partial files) from the HDD including slack space, unallocated space (Nolan, O'Sullivan, Branson, & Waits, 2005). Encase can also recover deleted emails from Personal Storage Table (PST file) in Microsoft Outlook.

Recovery options are limited once HDDs are wiped or data has been over written (Wright, Kleiman, & Sundhar, 2008). Gutmann (1996) published his research and wiped data can be recovered by *Magnetic Force Microscopy (MFM)* and *magnetic force Scanning Tunneling Microscopy (STM)* methods. Write (2009) concluded this is a misconception and his research demonstrated that using MFM to recover data from damaged HDD is achievable but "determine the prior value written to HDD was less successful than a coin toss". MFM is a imaging method measuring strength of residual magnetism (force gradient) on the platter surface with an optical interferometer or tunneling sensor (Gutmann, 1996).

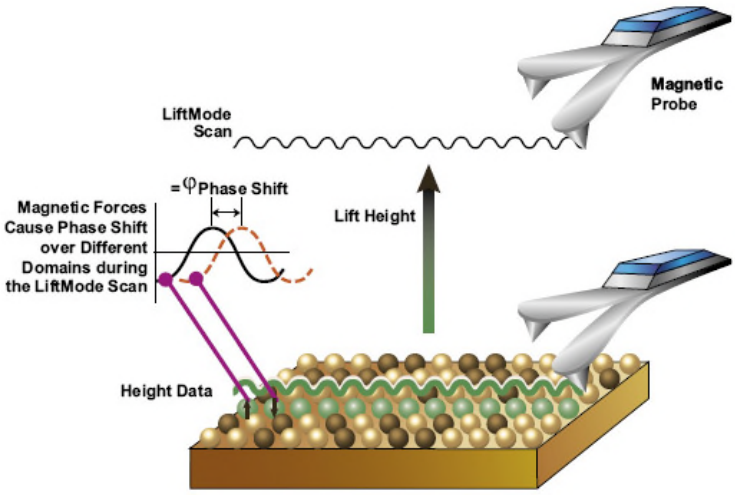

Figure 2.20: Illustration of MFM data recovery theory (Bruker, 2011)

STM use method similar to molding, typically covered with pure nickel on the platter surface, then peeling the thin film like layer and coat that with another layer of gold to prevent deterioration (Sawyer, 2006). Electrons on the film are tunneled through gap between a conducting tip mounted near the surface and *local density of states (LDOS)* across the surface are acquired (Chen, 1993). Although many enthusiastic hobbyists have built their own, this method is carried out at the atomic level, demands for exceedingly clean and stable surface, sharp tips, excellent vibration control, and sophisticated electronics (Mueller, 2006).

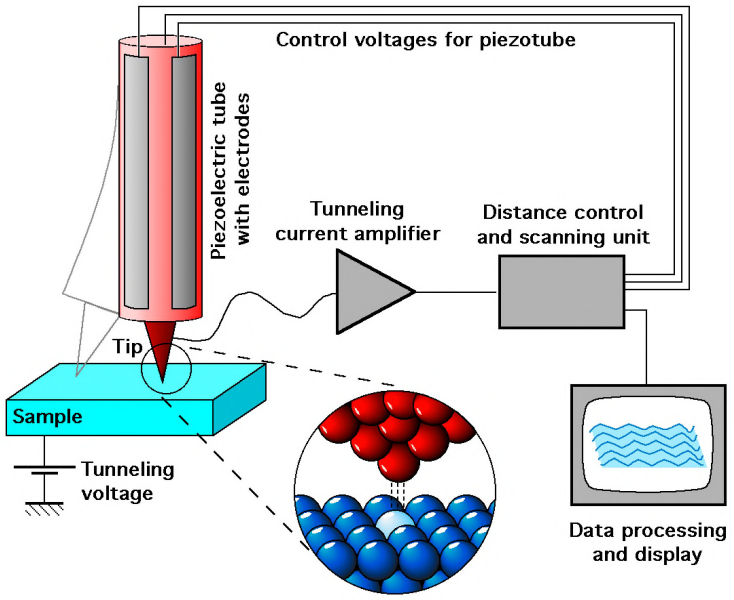

Figure 2.21: Schematic view of an STM (Schmid, 2009)

2.3.3 SSD Data Recovery

As discussed in Section 2.2, basic operation of SSD is significantly different from HDD. For example, “overwrite” concept prohibited in SSD operation. When a change has been made to a file, a new block space is assigned by the memory controller to copy the whole block with new information, then the original block will be marked as trash or available by the memory controller. This logical to physical mapping system is known as *Logical Block Addressing (LBA)* (Hu & Haas, 2010).

Regardless of different internal architecture, SSD was expected to act the same as HDD for the existing operating systems. *Flash Translation Layer (FTL)* was developed to enable this transparent compatibility. Minimum unit for write operation is one page, and minimum unit for erase operation is one block. When new data is entered, instead of replacing the changed page, a new page is saved and logical mapping is altered to point to the new location. The original location becomes obsolete and marked as “trash”. This is necessary because to write new data on written memory cells, it must be erased to write.

This inefficiency is multiplied with wear levelling and garbage collection, causing amount of physical write operation to increase exponentially in comparison to the logical amount intended, this phenomenon is known as *Write Amplification*. (Hu, Eleftheriou, Haas, Iliadis, & Pletka, 2009). This issue can only be mitigated and there is no permanent solution. Write amplication will not happen when cells are blank (Shimpi, 2009).

Factory default state of memory cells are blank, erasure is not required to write on a blank cell, and garbage collection won’t be executed till threshold of free space (blank) is breached. However wear levelling will slowly contaminate the blank cells and eventually write amplification will start to impact SSD performance. Some manufacturers use compression at SSD firmware level. This decreases the number of bits to store and has proven to be effective (Shimpi, 2009).

SSD has proven to demonstarate its ultra fast performance and steadily increasing its popularity despite higher cost per gigabyte when compared with HDD. However majority of consumers are not well informed about its limitations and technical glitches that accumilates over relatively short time of period. There are severe influences on digital forensics as well. Numerous research articles are

published in relation to SSD data recovery and security. In summary, unlike magnetic storage device "Data on SSD is really hard to erase AND tremendously hard to recover" (Bednar & Katos, 2011).

As discussed in Section 2.3.2, deleted data can be recovered from HDD, variety of commercial recovery software are available in wide range of price tags, and if the procedure is followed carefully the chance of getting the complete recovery is significantly high. Current forensic recovery rely strongly on this fact and even guidelines are formulated based on it.

Data recovery on SSD is delicate and can be more challenging especially for first responders. The moment data has been deleted, an artifact of deleted data on the memory cells are facing the risk of the automatic wiping algorithm known as garbage collection. Once cells are marked as obsolete (or trash), the memory controller will look for these cells so they can be written to thus leading to better oeverall performance. Once cells are cleaned, there is no way getting data recovered. In addition, there are two additional algorithm automatically wiping or in a way "overwriting" the deleted cells without any user intervention.

Trim is a new built-in command implemented into Windows 7 operating system or newer only to accommodate SSD's write delay by commnding SSD to efficiently retain blocks with least invalid pages (Hu, Eleftheriou, Haas, Iliadis, & Pletka, 2009). Unlike HDD, SSD does not overwrite and instead data is copied to new location without invalid pages and make more space available. The new location used could contained deleted files are wiped during this process.

Secondly wear levelling, both static and dynamic, can easily locate deleted space and attempt to evenly use all the memory cells. All these processes are carried out without a user's knowledge, even following the forensic acquisition guideline because there is no method available to stop them. Trim command can be blocked by a standard writeblocker because the command is sent from the operating system. It is worth mentioning that NAND flash memory's data retension period can be as short as few years, while HDD magnetic is proven to retain data for decades depending on the storage condition. Time is truly against data recovery process when it comes to SSD.

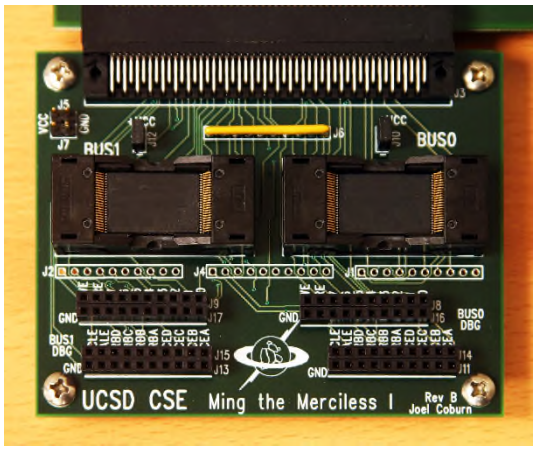

Figure 2.22: Custom FPGA-based flash testing hardware provides direct access to flash chips without interference from an FTL (Wei, Grupp, Spada, & Swanson, 2011, P.10)

The only and ultimate method of preventing such loss can be achieved by quickly remove SSD from power source, remove NAND flash memory chips from the PCB, then place them on a chip reader (e.g. *field programming gate array* or *FPGA*) to copy without the evidence destructing memory controller interferences (Winter, 2013). There are potentially three issues.

When data is copied diretly from the chips, extensive knowledge of file systems and *Redundant array of independent disks (RAID)* architecture is required to reconstruct the copied data. Due to the architectual difference, SSD data is not stored sequencially and rather random like a jigsaw puzzle. Fragments of traces and patterns are to be matched and data needs to be carefully reconstructed and therefore labour intensive and expensive.

To make the situation more complex, enterprises are learning to adapt a set of IT best practice, such as *Information Technology Service Management*, and becoming security concious especially on mobile computers. Although it is a good trend to be accepted in the wider community and confidential materials can be protected from careless accidents, when SSD data is encrypted it becomes almost impossible for any expert to recover the data. There is no way anyone can complete the puzzle without knowing what it should look like. Only if the encryption tool and password is provided, then there is a slight chance of recovering the data. However knowledge of the encryption algorithm is still required, in order to reconstruct the encrypted data.

Most manufacturers provide build-in encryption algorithms for improved security. If this hardware level encryption is in effect, the chance of data recovery

is highly unlikely even if the password is provided. This is due to an issue with lack of manufacturer support, and not being able to research their encryption algorithms. This is becoming a larger issue to the digital forensic community requirements. Not only manufacturer encryption algorithms but more and more vendor specific processes are added to newer SSD. Most of these are propietry to each vendor and details are not disclosed to public. Decades of HDD competition we have a handful of manufactures left dominating the current storage device market.

However SSD is new, and the level of entrance requirements is lower, especially for the memory controller. More than a hundred competitors across the globe are involved and making the situation more difficult for anyone to establish further standards specifically in relation to forensics and data recovery. Instead, current trends show the SSD manufactures are ignoring the forensic community. Lastly, and most important of all, removing chips by desoldering them from the PCB may damage the chip itself during the process.

This is known to be most intrusive method and unable to reverse the process once removed. The requirement of extensive knowledge and experience, precice hands-on surgical skills, a large amount of time and equipment, and potential risk of losing all makes this method not ideal and rather risky for the majority.

Data recovery on SSD is hard, but some researchers are concerned regarding the safety of deleted files, not being securely wiped. In the simplest terms, there is no secure way of manually wiping a data on SSD because data is not written sequencially, and a specific page can not be selectively sanitized (Freeman & Woodward, 2009).

Wei *et al.* (2011) has conducted research on reliability of SSD data erasure and concluded that when implemented correctly, existing built-in sanitization commands or software techniques are effective only for full-disk wipe, not for individual files. The researchers state the reason for failing to erase a single file is "because FTL complexity makes it difficult to reliable access a particular physical storage location" (Wei, Grupp, Spada, & Swanson, 2011). The authors proposed *scrubbing* method to securely erase data at individual file-level by programming pages and turning remaining 1s to 0s (Wei, Grupp, Spada, & Swanson, 2011).

The result was acceptable with SLC, but MLC suffered severe corruption because a single memory cell is shared by 2 bits which belong to other page (Wei, Grupp, Spada, & Swanson, 2011). Overall, the research shows there is a potential to securely erase data at file level, however it still lacks ability to verify the secure deletion after scrubbing an individual page. Erasure without verification is a great concern for the forensic community. Further research and development is required in this field.

2.4 DIGITAL FORENSICS

When electronic devices are used or involved in an incident or crime, a special examination technique called *digital forensics* is used for the recovery and investigation (Carrier, 2011). This research mainly focuses on the collection phase.

According to the The Kipling method (5W1H) is often found useful during this process (MediaSmarts, 2013). Otherwise if the hashes do not match or are irreproducible then the acquired data is considered tampered and becomes an unreliable source (Travis & Rau, 2004).

In order to achieve such rigorous requirements, the Department of Justice (DOJ) National Institute of Justice (NIJ) of the United States published the Forensic Examination of Digital Evidence: A Guide for Law Enforcement in 2004 (Travis & Rau, 2004). This guideline suggests to develop, test, and document technical digital forensic procedures for reproducible outcomes, and the establishment of standard procedures ensures identical data processing (Clausing, 2009). A standard procedure covers from the search warrant, taking scene photos, documentation, secure storage container and transportation. This research covers data acquisition and procedures.

2.4.1 Forensic Imaging

The term forensic imaging is often used as a synonym for the acquisition, because no matter how the evidence is seized, analysts shouold create an identical duplicate copy and save as an evidence image format for examination. The best digital forensic practice requires for minimize (or ideally “eliminate”) any modification to the original media in order to retain the integrity and reproducibility (Britz, 2009). The U.S. NIST uses Computer Forensics Tool Test (CFTT) program to ensure the consistent reliability and accuracy of computer

forensic tools (Lyle, 2003). The CFTT establishes a methodology for testing the disk imaging tools and results are used by users to make informed decisions (Lyle, 2003). Below are a list of major elements tested on imaging tools by the CFTT.

- Features (copy, image, verify);
- Target media (ATA, ASPI, legacy Bios, Bios to IDE/SCSI);
- Size of results (source size compared with destination size);
- Errors (source read, destination write, image read/write/change);
- Target format (Disk, FAT12/16/32, NT, Ext2); and
- Remote access (abailability).

The Write-blocker is a technology duplicates a source drive without corruption or alteration. It is available in both hardware and software. Tableau, WiebeTech, and MyKey are the leading hardware write-blocker (HWB) manufactures. Modern disk imaging tools, such as Tableau TD2, can acquire duplicate copies into two storage media at once (also known as "twinning") without corrupting the original source (Marshall, 2009). The write-blockers can not only prevent alteration to the source drive, but also regulate the speed to read at slower speeds for troublesome drives and stop malicious programs or physical damages to causing further damage to the source (Jang, Koh, & Choi, 2012).

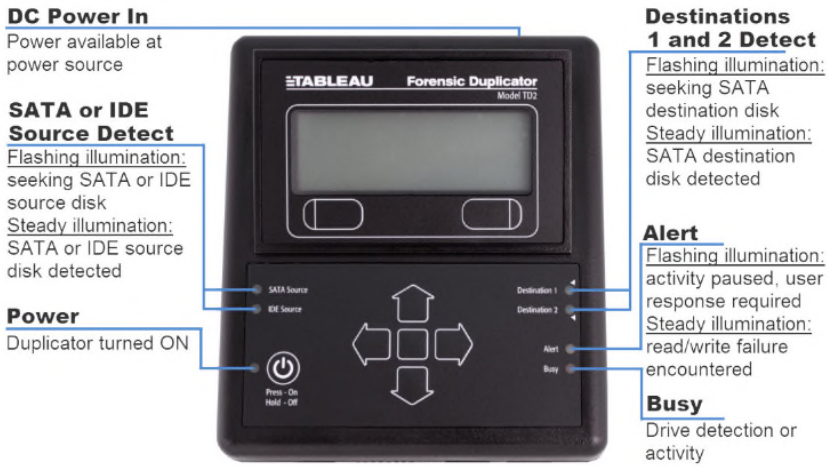

Figure 2.23: Picture of Tableau TD2 write-block duplicator (Guidance, 2008, p.6)

There are several image formats made available for forensic imaging shown in table 2.3. Each format has slightly different advantages but Encase image file format (E01) is often considered the de facto standard for computer forensic

analysis. Shown in Table 2.3, E01 formart is widely accepted by various commercial forensic tools where other formats, other than raw image format (.dd), are limited to specific tools and significantly lacks compatibility.

Table 2.3: Evidence image formats supported by each forensic tool (Mueller, 1999)

Format Tools	AFF	EnCase	Expert Witness	ProDiscover	Raw	SMART Comp
AFFlib	√	√	√		√	
EnCase		√	√		√	
FTX		√	√		√	√
ProDiscover				√	√	
Sleuth Kit	√	√	√		√	
SMART		√	√		√	√
X-Ways		√	√		√	

E01 format design is heavily based on its predesessor Expert Witness Compression Format (EWF) developed by ASR data (Jang, Koh, & Choi, 2012). Garfinkel, et al (2009) describes that E01 image is prefixed with a header, followed by a physical bitstream of an acquired disk which is interlaced with CRCs for every block of 32KB (64 sectors), and closed with a footer containing an MD5 hash for the entire bitstream (Garfinkel, Malan, Dubec, Stevens, & Pham, 2006).

The header consists of date and time of acquision, examinar's name, notes, optinal password, and own CRC is used to conclude the header. The header is also refered as a "case info" (Figure 2.24). Advanced features of E01 format are the ability to compress and search, but it is unable to provide a digital signature for authentification and encryption for confidentiality (Garfinkel, 2010). In addition, due to its propietry format architecture not all the details of E01 are fully disclosed.

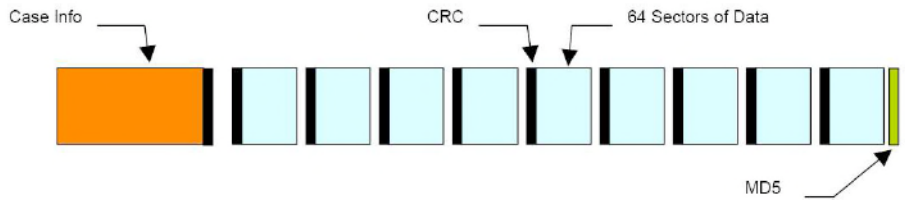

Figure 2.24: E01 file structure (Garfinkel, Malan, Dubec, Stevens, & Pham, 2006, p.15)

Advanced Forensic Format (AFF) has developed to solve all the above limitations with E01. AFF is an open source forensic image format, supports multi-platform, public/private key encryption, and stores large data without splitting into multiple

files. As Garfinkel, et al (2009) states in their research, AFF is a flexible format that can be used for a variety of tasks and overcome limitations E01 suffers. Therefore three types of forensic image format should be considered for this research. Raw image for simple, fast handling, E01 for the best compatibility and widely accepted, and AFF for advanced flexibility removes all hindrances.

Forensic images must be verified by using SHA-1 or MD5 hash algorithm to ensure data are duplicated at sector level. This verification will not only guarantee its completeness but also provides security towards authenticity through out the investigation process.

There have been numerous attempts to develop a specialty guideline (or so called "best practice" guide) for digital forensics but none has been accepted universally. Best practice is a method or technique necessary to maintain quality and also provides standard procedures to all experts involved in an investigation (Bogan & English, 1994). The Association of Chief Police Officers (ACPO) Good Practice Guide for Computer based Electronic Evidence or NIST 800-86 Guide to Integrating Forensic Techniques into Incident Response are commonly referenced as one of best practice for digital forensics. Availability of widely accepted forensic processes allows any forensic analyst to reproduce the same output. Being able to reproduce the same result by following the same processes is a critical part of digital forensic analysis.

2.4.2 Storage Device Preservation

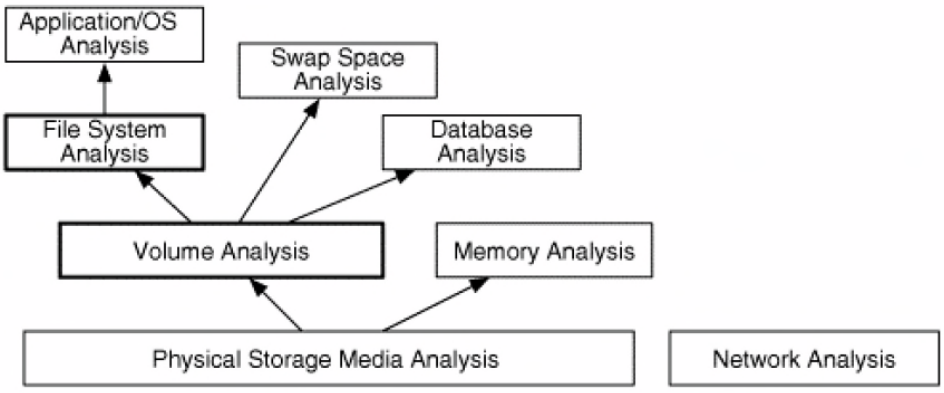

Figure 2.25: Process of analysing data at the physical level to the application level (Carrier, 2011, p.12)

The general digital forensic approach of acquiring a storage device, a HDD in particular, is to copy sector-by-sector or bit-by-bit known as a *bit-stream image* (Kent, Checalier, Grance, & Dang, 2006). An Analyst can obtain a perfect

duplicate of the source media to a suitable removable media (Britz, 2009). Carrier (2005) describes a HDD consisting of four main layers; physical (or disk), volume, file system and application. Data are lost at each layer of abstraction (or higher layer), therefore bit-stream image of the lowest, disk layer, should be acquired (Carrier, 2011). To clarify further, if a HDD was acquired at volume layer level, that will allow recovery of deleted files in each partitions but is unable to analyse sectors not allocated to partitions, also known as *unallocated space* (Carrier, 2011). As shown in Figure 2.26, an unallocated space is a space in between partition or a space not being used by any partition on the HDD, and data can be hidden in this area.

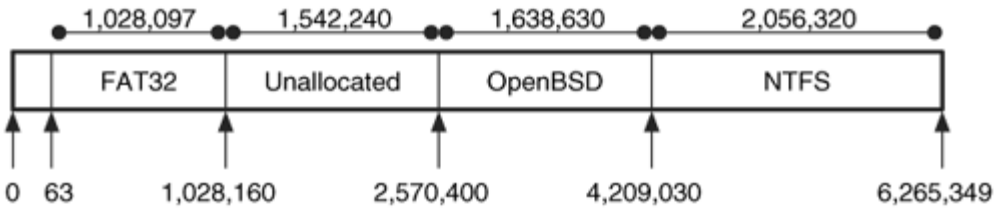

Figure 2.26: Layout of the example disk image (Carrier, 2011, p.78)

In addition, if an acquisition was processed at the file system layer, only allocated files can be analysed and areas such as unused space (also known as *slack space,* often containing residual data), deleted files, temporal data, hidden data within partition or file system structure will not be available (Kent, Checalier, Grance, & Dang, 2006).

There are two types of acquisition, live and physical (or traditional), but both aim for the same goals; reduce the amount of modification made to the original source, and the same copy of the original can be recreated if necessary (Kent, Checalier, Grance, & Dang, 2006). The live acquisition is used where the system is currently in use and volatile data such as temporal cache files and random access memory needs to be acquired. In corporate environment or covert operation often prohibits the analyst to switch off the power and data must be copied live from the source media.

Alternatively traditional digital forensic acquisition deals with physical drive itself. The best practice suggests to use a write-blocker (software or hardware) to prevent modification (or overwritten) to be made on the source media (Britz, 2009). Once an identical copy is saved on destination media, a hash value such as *Message Diget 5* (MD5) or *Secure Hash Algorithm* (typically SHA-

1) is calculated to retain integrity. If a single bit of data is changed then the hashes will change significantly. Therefore if the hash value does not match then the anaylst will know the data has been modified (Carrier, 2011). The hash value for the original source media is calculated prior to the imaging. A copied media's hash value is calculated after the imaging (Kent, Checalier, Grance, & Dang, 2006). In September 2012, NIST has announced that Federal agencies in the United States should stop using SHA-1 for generating digital signatures (hash values) and use *SHA-256* at a minimum for any application requires hash values (Barker, Barker, Burr, Polk, & Smid, 2012). SHA-256 is one of four hash functions made available in *SHA-2* (NIST, 2002).

Another reason the bit stream imaging is recommended in the standard digital forensic procedure is that its ability to preserve *file times*. Three types of below file times are known as *MAC times*. Analysts often use the file times to reconstruct a timeline for system activities. There are differences the way a time stamp is handled amongst operating systems.

For example, in some UNIX (widely used multi-user operating systems, such as BSD and SunOS) last accessed time for executable files are not updates when they are run (Kent, Checalier, Grance, & Dang, 2006). Therefore analysts must be aware of not all acquisition method can preserve file MAC times, and different operating systems iterpret MAC times disssimilarly.

Table 2.4: Descriptions for the file modification, access and creation time stamps (Kent, Checalier, Grance, & Dang, 2006)

Modified time	Describes when the file was most recently written or changed.
Access time	Identifies the time file was most recently read, or opened.
Creation time	Identifies when the file was created (or copied to the new system)

2.4.3 Digital Forensic Guidelines

It is apparent existing digital forensic processes are challenged by SSD requirements. Ultimately a court of law demands for assurance that forensic tools and procedures are accurate and produce reliable results. This is where a guideline has a place and plays a major role. American Standards of Testing and Materials (ASTM) defines that "Standard test method is the way a test is performed. Standard practice is a sequence of operations that, unlike a test, does not produce a result. Standard guides provide an organized collection of information or series

of options that does not recommend a specific course of action" (Ballou & Gilliland, 2011).

Similarly Trček et al. (2010) concluded their research that by providing a new solution based on methodological and procedural approaches in digital forensics, it can provide more reliable data and "significantly ease forensic investigations", and "it is worth to compare frameworks to establish a best practice guide".

Number of reputable institutions have published guidelines for digital forensics. Those guidelines are occasionally revised to accommodate newer technologies (e.g. smartphones and wireless network forensics). Six de facto guidelines are methodically reviewed and compared for their features. Note that only data preservation, acquisition, and recovery processes are mainly focused on in this review.

2.4.3.1 Scientific Working Group on Digital Evidence (SWGDE) - Best Practices for Computer Forensics

The latest best practice guideline for computer forensics is published by Scientific Working Group on Digital Evidence (SWGDE) in February 2013. SWGDE is a sub-division of Scientific Working Groups (SWGs), which consists of lead scientists across the globe within the focus field, and established to improve discipline practices and develop consensus standards (Sammons, 2012).

Their guideline "Best Practices for Computer Forensics" is designed to provide basic computer forensics practices for forensic examiners and first responders. Information is thoroughly composed with sufficient details in six pages. It is documented sequentially with bullet points. Critical technical terms are well explained without breaking the step by step flow. No graphical representation is used in the guideline and Figure 2.27 should assist visualise work flow where a typical data storage device requires investigation.

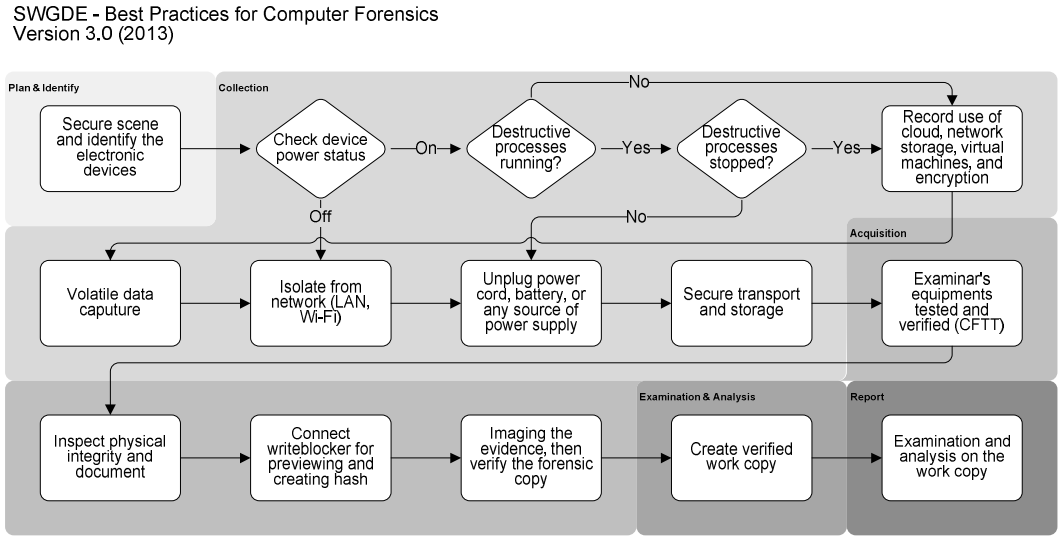

Figure 2.27: Process flow chart for SWGDE best practice for computer forensics

The flow and involved processes are simplified and focusing especially on preservation and acquisition phases. Five main forensic phases are colour coded to visually simplify the processes involved in each phase. The basic forensic flow of plan, identify, secure, collect, acquire, examine, analyse, and report are all available and sufficient details are provided. However neither NAND flash nor SSD is mentioned in the guideline and hence following this guideline will significantly increases the risk of incomplete data recovery.

2.4.3.2 Association of Chief Police Officers (ACPO) - Good Practice Guide for Computer-Based Electronic

Association of Chief Police Officers (ACPO) is a private company in England, providing a forum, policies and guidelines for chief officers as well as local forces across England since 1948. (ACPO, 2010)

ACPO offers the second latest guidelines amongst the others, and two most recent editions of "Good Practice Guide for Computer-Based Electronic", fifth version (issued March 2012) and fourth version (issued December 2007), are reviewed for this research. The guidelines are covering a wider range of aspects at a high level. Basic forensic practice flow is maintained but generally providing information broadly and maybe suitable for upper management or higher level. It is near impossible to draw a flow chart from fifth version guideline. ACPO has made the guideline even broader to encompass emerging technologies and diversity of the cyber security incidents (Wilkinson, 2012).

ACPO states that majority of computer based evidences are retrieved from standalone or networked computers such as desktops and laptops (Williams,

2007). The work flow and processes suggested in the guideline is sketched in Figure 2.28. The fifth version did not provide sufficient details and therefore the majority of processes are derived from fourth version.

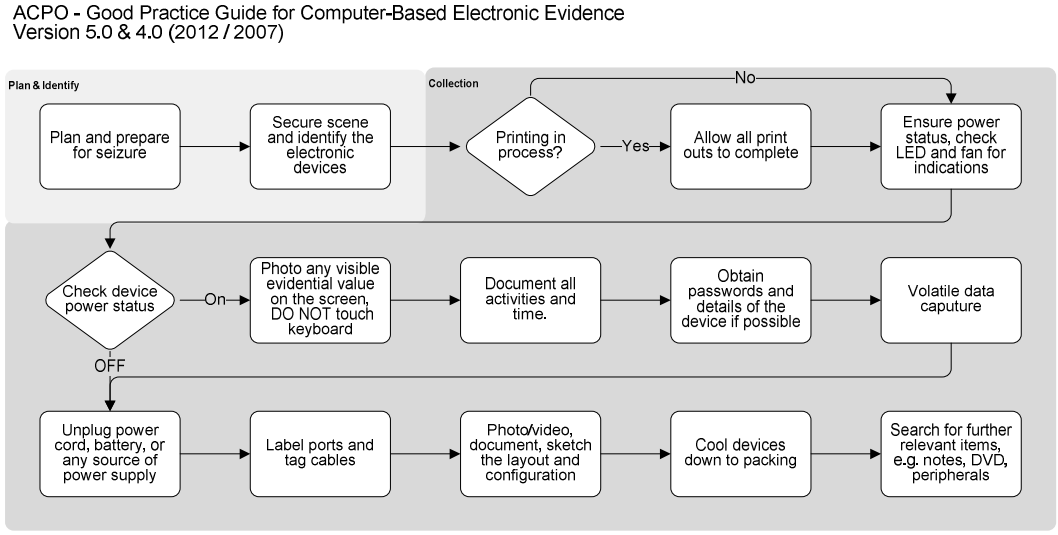

Figure 2.28: Process flow chart for ACPO good practice guide for computer-based electronic evidence

Similarly a flow chart was available for digital CCTV forensics but not for computer forensics. In comparison to SWGDE flow chart, it is obvious that the ACPO guideline provides details limited to two phases. There is no information made available for after collection and the guideline is making various suggestions and considerations, not instructions. Positive aspects of the ACPO guideline are that this is the only document mentioning SSD's name and instructions. Potential data preservation issues with SSD technology are not discussed at all.

2.4.3.3 National Institute of Justice (NIJ) - Electronic Crime Scene Investigation: A Guide for First Responders, Second Edition

National Institute of Justice (NIJ) is a governmental research organisation (agency) of the United States Department of Justice, originally founded since 1968 (NCJRS, 1998). NIJ specialises their research in advanced technology for criminal justice such as forensics. "Electronic Crime Scene Investigation Guide" was first published in 2001; the second edition "Electronic Crime Scene Investigation: A Guide for First Responders" was published in 2008 which was reviewed for this research. As the title suggests, this guideline is composed for first responders and written in plain language as well as step by step instructions. For this reason, the guideline introduces detailed tools and techniques for collection (seizure) but

nothing else, which perfectly fits for the role of first responders. See Figure 2.29 for the NIJ work flow which was derived from NIJ's collecting digital evidence flow chart.

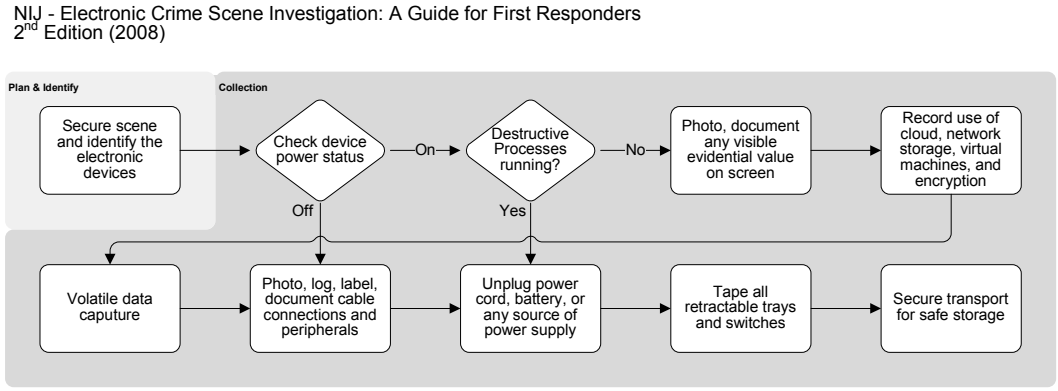

Figure 2.29**: Process flow chart for NIJ electronic crime scene investigation – a guide for first responder**

Each step is clearly described in chronological order, but the focus is limited to the collection phase, and finishing with secure transportation and storage. In another publication from NIJ, "Forensic Examination of Digital Evidence: A Guide for Law Enforcement", which targets more on law enforcement agents not only for collection but examination and analysis. However due to the broader scope coverage, details are not available, and step by step nor flow charts are not included. It was published in 2004 and SSD or NAND flash are not discussed at all. Case examples are included in the appendix and can be useful for a research experiment when the performance of different guidelines is compared.

2.4.3.4 National Institute of Standards and Technology (NIST) - Guide to Integrating Forensic Techniques into Incident Response

National Institute of Standards and Technology (NIST) is a non-regulatory agency of the United States Department of Commerce, promoting the country by establishing lead in measurements, standards and technologies since 1901 (NIST, 2008).

"Guide to Integrating Forensic Techniques into Incident Response" is written from technical view, not law enforcement. It is a great technical reference point for incident response teams or computer forensic analysts, providing detailed information for identification of evidence, preservation method without losing data integrity, and how forensic equipment needs to be maintained and tested at all time.

However this publication is more close to a literature and it should serve greatly as a source of information for consideration when forensically acceptable technical approach is in question. It provides what to consider and the reason why, but not how. Figure 2.30 shows brief flow chart based on the NIST guideline where a typical computer with data storage device requires an investigation.

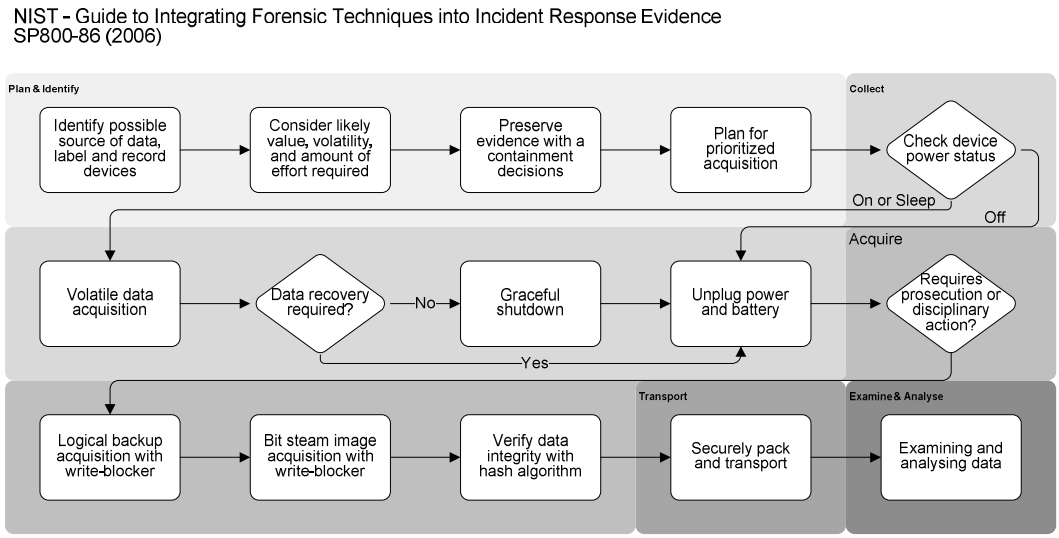

Figure 2.30**: Process flow chart for NIST – guide to integrating forensic techniques into incident response evidence**

Overall coverage was sufficient but it significantly lacks clear instructions and rather introduces suggestions. Although a variety of technologies are discussed, SSD or NAND is not included. It has been seven years since the first publishing but there is no sign of a revised version was found from the NIST website.

2.4.3.5 SANS (SysAdmin, Audit, Networking, and Security) Institute – Forensic Plan Guide

The SANS Institute was originally formed as technical training/conference purpose in 1989, belonging to the for-profit company Escal Institute of Advanced Technologies and specialises in research archive, training and professional certification (SANS, 2013).

"Forensic Plan Guide" was published in 2006, attaching "Forensic Cook Book" to cover further practical forensic guideline with sample tools and case examples. This is by far the most comprehensive computer forensic guideline, covers all the relevant technical fundamentals yet provides clear guidance of complete forensic procedures with sample freeware Linux forensic tools and examples. High-level work flow in relation to data collection and recovery is shown in Figure 2.31.

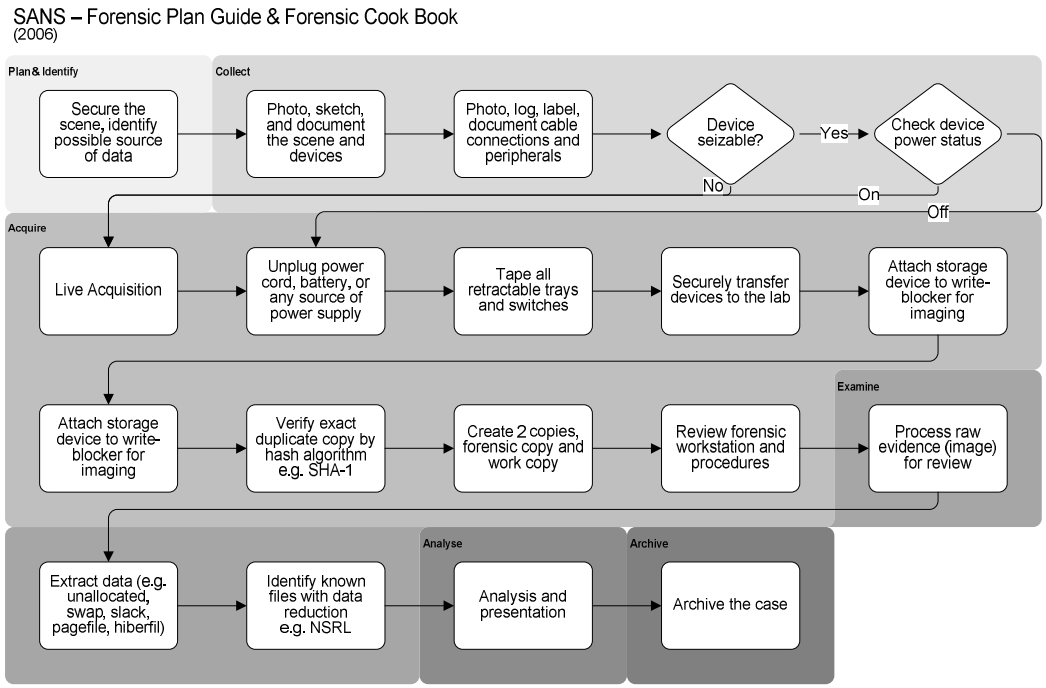

Figure 2.31**: Process flow chart for SANS – forensic plan guide and forensic cook book**

Strictly speaking, this guideline is not composed ready for everyone. Expecting all first responders to flawlessly follow this guideline is ideal but unrealistic. It is hard to say if SANS guideline is better than NIST's. While NIST's provide wider options and suggestions for considerations, SANS is more directive and practical. When those two guides are compared, there is no significant difference between them and both cover sufficient detail from the initial investigation to the end of the forensic cycle.

It is disappointing to see that both are equally out dated and not covering the known issues with SSD and NAND flash storage devices. Obviously further research is required to inform all personnel involved in computer forensics that a new generation of storage devices is increasing popularity in the market, numerous experts in this field are alerting others to the risks involved, and guidelines must be updated accordingly.

2.4.3.6 Computer Emergency Response Team (CERT) - First Responders Guide to Computer Forensics

Computer Emergency Response Team (CERT) is a trademark for a team responsible for security incidents at Carnegie Mellon University and also known as Computer Security Incident Response Team (CSIRT), formed in response to a worm epidemic which paralysed IBM VNET in 1988 (SEI, 2010).

CERT's forensic guideline "First Responders Guide to Computer Forensics" was published in 2005, the oldest in this review, and has not been revised or updated since. It was reviewed because the organisation is coming from a strong academic research background similar to SANS and NIJ. The most interesting part of this guide was their definition of *persistent data*, and focusing on loss or contamination of these data due to improper handling. Persistent data is "retained and remains unchanged after the computer has been powered off" (King & Vidas, 2011). However the contents are dissatisfactory, lacking the forensic work flow and basically describing key components, terminologies, and introducing various relevant forensic tools. Figure 2.32 is shows a work flow based on CERT's guideline.

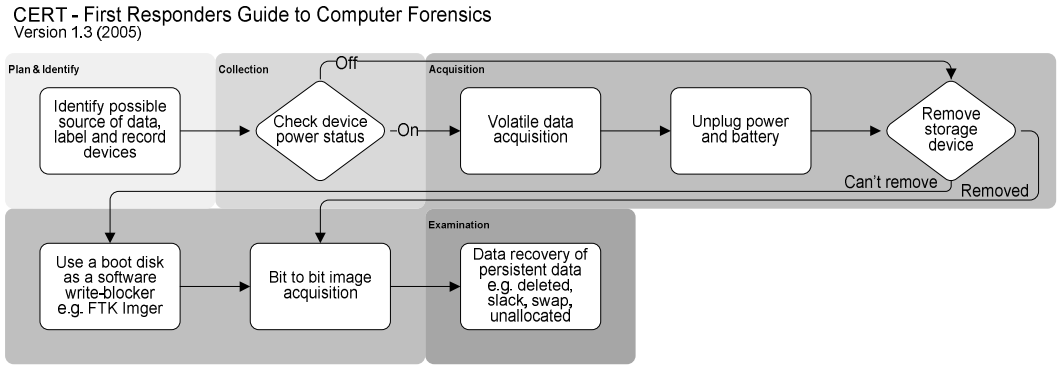

Figure 2.32**: Process flow chart for CERT – first responders guide to computer forensics**

Due to the lack of adequate detailed processes, it is apparent the flow chart is overly simple and insufficient. Suggested tools and techniques are also well covered in other guidelines. This guide maybe suitable for entrance level first responders, but the contents are old and desperately require major refurbishment for practical use.

2.5 CONCLUSION

SSD manufactures are succeeding to mass produce less expensive NAND chips with increasing multi-bits per cell technology, consumers are rapidly converting from HDD to SSD without learning the potential downsides. The increase in the consumer usage means digital forensic analysts are exposed to the new challenge. Existing guidelines are still lacking protocols to prevent loss of data when SSD is handled same as traditional HDD. A new guideline to mitigate risks involved with handling SSD must be researched.

The literature review in Chapter 2 identified HDD and SSD are two distinguishingly unique storage device developed for the same purpose. Although it appears the same to the end users, distinct architectural variance delivers new challenges to the digital forensic investigation. None of data preservation and recovery methods are applicable to SSD, and residual data are automatially wiped without human intervention. Even extensive approaches are taken to recover data from SSD, data is often encrypted and compressed with a vendor proprietary algorithm. The vendor proprietary algorithms causes two major issues. Lack of standards and less successful data recovery.

Base on above factors, a number of issues are found and the situation will worsen as SSD continues to grow its market share. The guidelines reviewed in Section 2.4.3 must incorporate processes to accomodate events where SSD is present. Chapter 3 will discuss research methodology, research design, data collection and analysis in detail. Similar researches will be studied to identify a research opportunity and establish a hypothesis.

Chapter 3

Research Methodology

3.0 INTRODUCTION

Research methodology is an analysis of how information is captured and processed for a field of study for a specific objective. Chapter 2 identified issues based on literature reviews in modern data storage architectures and generic digital forensic operations. SSD has become affordable and manufacturers are investing into the technology as a substitute to HDD. Digital forensics is never ahead of modern technology and research and development work takes significant part of forensic analyst's time.

The main forensic issues with SSD are inability to duplicate and recover data. Data on a SSD can be copied to other media but when a forensic image is made, hash verification is difficult. Mismatching of hashes generally indicates the two copies are not a duplicate copy, the integrity of the copy is lost, and therefore admissibility in court of law becomes highly doubtful.

SSD data recovery is challenging because of built-in processes constantly making deleted space sanitised and making the space ready for new data. These processes are fully automated and nothing can stop them unless NAND chips are extracted from the circuit board. These processes are also the cause for mismatching hash values.

Evidence collection is typically handled by trained first responders. Training are based on well established guidelines which were critically reviewed in Section 2.4.3. The literature review in Chapter 2 indicates the guidelines are a critical part of the digital forensic lifecycle, but the distinctive differences between HDD and SSD architecture are yet to be critically considered.

Chapter 3 will discuss how SSD capability on current guidelines can be systematically and theoretically analysed for research. Section 3.1 will review literature from similar studies for methodologies, benchmarks, and further refinement to the research scope. Research questions and hypotheses are defined in Section 3.2, information flow will be presented in a data map. Research requirements such as data collection criteria, methods, processing of collected

data and analysis techniques are presented in Section 3.3. Research limitations are discussed in Section 3.4 and then a conclusion is made in Section 3.5.

3.1 REVIEW OF SIMILAR STUDY

In this section, four similar research reports relevant to SSD data recovery and digital forensic guideline analysis are reviewed. Section 3.1.1 examines critical analysis on existing guidelines and SSD. Section 3.1.2 and 3.1.3 examine studies conducted to experiment with secure data deletion and success rates of data recovery. Section 3.1.4 examines how SSD and HDD behaviour are tested and what recommendations are made from the results.

3.1.1 SSD Forensic Guideline

Bednar and Katos (2011) recognised SSD is a completely new technology, which imitates HDD behaviour but has serious consequences for digital forensic investigation. The evidence destruction behaviour of SSD is identified as the most concerning issue and the authors suspect the cause of the issues are probably an artificial problem, such as lack of widespread understanding and maturity of SSD technology. The latest ACPO guideline and potential of live acquisition is also discussed. The effect of garbage collection and TRIM processes are noted with concern. Furthermore the authors discuss that these processes are defeating the purpose of write-blocking, making it impossible to create a traditional forensic duplicate copy.

The authors conclude that although the ACPO guideline is carefully documented and valuable resource, the information cannot be applied and is not suitable for handling SSD collection or even not feasible to develop a hybrid guideline incorporating both HDD and SSD. This is due to the fundamental technology being completely different. In addition, although NAND flash chip extraction data recovery methods were not practically established back then, the authors commented that such extreme acquisition involving physical tampering will cause a non-trivial exercise for validation, certification, and forensic compliance (Bednar & Katos, 2011).

Although garbage collection is a firmware built-in process, TRIM is a command only available in modern operating systems, such as Windows 7, and it can be manually disabled. As Bednar and Katos (2011) identified TRIM as a

variable factor which potentially influences live forensic processes. Further testing is required to establishing new guidelines. The study was qualitative research, but if success rate of live forensic processes can be measured and compared then the result should verify the authors' conclusion. If the result supports the conclusion, that suggests there is no guideline currently available to handle SSD at crime scenes.

3.1.2 SSD Data Retention Analysis

Empirical analysis was used by King and Vidas (2011). 15 different SSD and one HDD were involved in this research to test the applicability of the data loss using three case scenarios. The result shows no data can be recovered when TRIM is enabled, and when TRIM is disabled the recovery rate improves significantly but differs depending on the manufacturer. The authors suggested that while TRIM is enabled, traditional data recovery is no longer viable option for investigators.

As discussed in Section 2.2, the issues with SSD are firstly TRIM and garbage collection executed by the flash controller that result in the equivalent of a sanitized disk (Kissel, Skolochenko, & Li, 2006). Secondly wear levelling on SSD disallows access to remnants of data through FTL (Bell & Boddington, 2010).

The research shows that without TRIM enabled averages near 100%, the rate of recovery vary depending on SSD manufacture and with TRIM enabled (both operating system and SSD), no data was recoverable (King & Vidas, 2011). The authors conclude that based on their experiment, forensic investigators should identify the use of SSD and the operating system before pulling the power plug, because if TRIM is enabled the chance of data recovery from the SSD is unlikely and volatile data analysis may assist with the investigation better.

The results are also supporting Bednar and Katos' (2011) research discussed in Section 3.1.1. Although details of how the data recovery was performed were not clearly discussed, it only states "files were recovered from each drive after deletion" (King and Vidas, 2011). This raises questions if various live and dead forensic acquisitions were used, would the different acquisitions produce similar results?

A similar study comparing SSD data retention and TRIM function was published in end of 2013 by Nisbet, Lawrence and Ruff (2013). Multiple

operating systems were involved in their study to examine differences between file systems and data recovery performances. Nisbet et al. (2013) adapted data collection methods from King and Vidas's (2011) research. This was data retention across TRIM enabled file systems and the use of the TRIM command as an alternative method to sanitise SSD. The results showed TRIM can effectively purge deleted data within minutes. When data is purged on SSD, no residual data can be recovered. The authors also identified how the TRIM command is executed differs depending on file systems, especially the Ext4 file system used in Lunux Kernel version 3.6.33 or later, which executes the TRIM command in batches (Nisbet, Lawrence & Ruff, 2013). Batch processed commands are less instant and therefore the chance of recovery is slightly higher than others.

In comparison, Nisbet et al. (2013) also tested the same drives without TRIM enabled. The authors report that when the drive usage is high without TRIM enabled, a more aggressive garbage collection process was monitored. In other words, disabling TRIM may not sufficiently suppress the automatic destruction process, but it depends on SSD storage space usage as well. Various conditions seem to influence chance of data recovery with SSD.

3.1.3 SSD Secure Deletion Efficacy

Freeman and Woodward (2009) demonstrated none of files could be recovered from securely deleted SSDs. Some files can be carved but none were loadable. The main purpose of this research was to measure the efficacy of secure deletion on SSD, but methodologies used to sanitize SSD and measure data recovery performance are desirable.

In the experiment, the same selected file types are deleted from each test SSD, and then carved to discover if any files were recoverable from the device (Freeman and Woodward, 2009). File type, size, and status of recovery and loadable are recorded in a table.

As the authors mentioned in the conclusion, TRIM was not considered in this study and secure data erasure on the SSD was the main objective. As an extension to this research, various acquisition methods based on case scenarios can be performed and use one of the carving tools to measure the rate of data recovery.

In the recent study, Nisbet et al. (2013) presented the result of using TRIM as an alternative secure erase method. The authors concluded that the erasure speed is fast and effective, but minimal data remain untouched. Therefore if SSD supports ATA Secure Erase, the study recommended not using TRIM as an alternative.

3.1.4 SSD Data Recovery Guideline

Bell and Boddington (2010) have conducted similar experiment prior to King and Vidas (2011), but only data recovery on formatted SSD were tested and not deletion. The research shows only 0.34% of recoverable files survived (1,090 of 316,666) and none could be recovered in the complete original form. Garbage collection deleted metadata supporting the existence of 96.66% of files within a few minutes during the recovery (Bell & Boddington, 2010).

Because these destructive operations are carried out within SSD (between NAND flash and flash controller), attaching a write-blocker between the computer and SSD did not prevent the internal processes. The authors state that due to the SSD's own decisions in the absence of computer instructions, "it is prudent and potentially reckless to rely on existing evidence collection processes and procedures" and for those reasons "the golden age for forensic recovery and analysis of deleted data may now be ending" (Bell & Boddington, 2010).

Bell and Boddington (2010) suggested data stored on NAND flash base media should be considered as "grey area" in terms of data recovery and legal verification, and further extensive research is required. The authors stated data preservation on SSD cannot be guaranteed once data is marked as deleted, although none of live acquisition methods were tested. Both software and hardware write-blockers are suggested as not capable of preventing internal process such as garbage collection. Again only USB write-blocker was involved in the research and a software write-blocker was not tested. It would be informative to test if certain combination of acquisition processes with a software write-blocker provides any different result.

As the authors mention in the study, there is a possibility that a combination of certain drives, firmware, and acquisition process may facilitate or prevent unintended permanent data loss (Bell and Boddington, 2010).

3.2 RESEARCH DESIGN

Section 3.1 critically examined related research reports and identified number of research questions. Section 3.2 will process the information gathered from the literature review in Chapter 2 and Section 3.1, and formulate a research design. Section 3.2.1 will discuss the findings in Section 3.1 and derive an overall objective for this research. The selected research questions are discussed in Section 3.2.2, and followed by research hypothesis in Section 3.2.3. Section 3.2.4 will discuss the research phases to clarify the work flow, and finished with visualising the research design in a data mapping in Section 3.2.5.

3.2.1 Review of Similar Studies

A critical review of relevant studies was completed in Section 3.1 and a summary of findings are discussed to set the direction of this research in this section. Overall, the results from the similar studies demonstrated that an acquired SSD forensic copy cannot be verified with the current acquisition methods. This is caused by built-in firmware operations designed to provide better performance without considering any impact on data recovery or preservation. This leads the authors to question the existing guidelines and suggest that the presence of SSDs make it difficult to develop a hybrid guideline which incorporates HDD as well as SSD. However none the research has critically analysed the guidelines in relation to SSD acquisition capability, nor tested potential combinations of acquisition processes to verify their theories.

Although SSD firmware based processes are difficult to intervene without NAND chip extraction, the literature review identified that the TRIM command can be disabled. TRIM is only effective on later operating systems and it's influences can be tested by simple live acquisition comparison. Effectiveness of live acquisition can also be tested when same data set and recovery methods are used in controlled environments. The myth of utilising combinations of various SSD, write-blocker, forensic imaging tools, can also be tested using test case scenarios. Measurement of results can be quantified and analysed by adapting data recovery rate research methods used in Bell and Boddington (2010). This research will fill the gaps identified in Section 3.1 and be investigated if the current guidelines are truly incapable of incorporating the SSD forensics.

3.2.2 Research Questions

A summary from the review of similar studies in Section 3.1 and in-depth analysis of the major guidelines in Section 2.5.3 is given in Figure 3.1. The analysis delivers a visual guide to what is currently known and a roadmap for research.

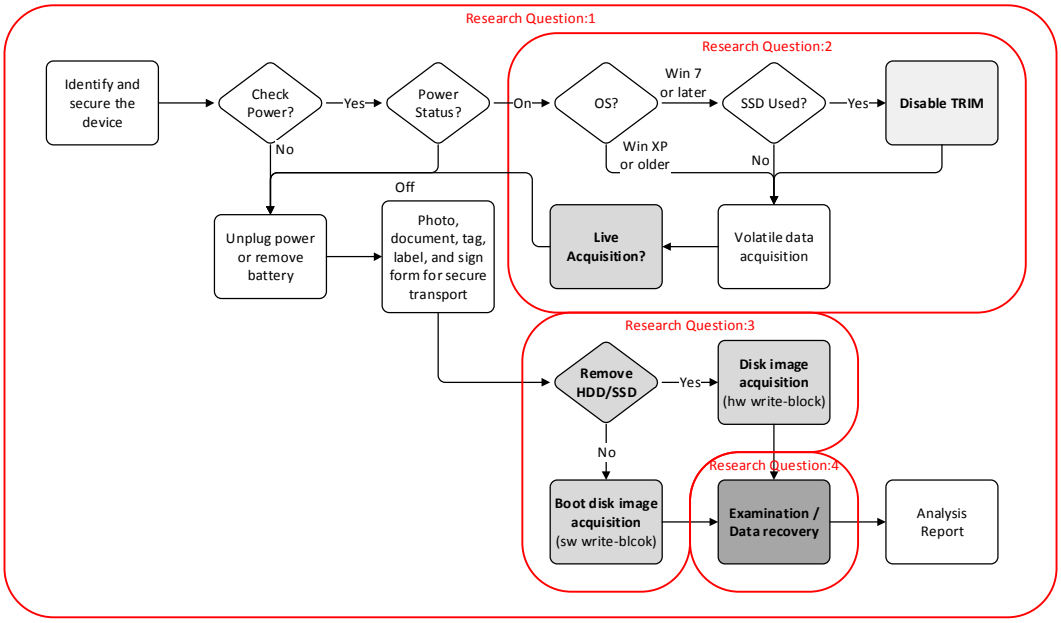

Figure 3.1: Compound acquisition chart with identified sub-question processes - illustrating Potential Area of Data Recovery Improvement

The research question and sub-questions are consequently formulated based on the selected acquisition processes. The research question for the thesis is thus:

Is it feasible to adapt existing guidelines for SSD forensics?

Consequently sub-questions are directly associated with the selected processes and these are listed as follows.

Sub-question 1: What is forensically acceptable SSD acquisition?

Sub-question 2: What changes can be observed if TRIM is disabled prior to live acquisition?

Sub-question 3: What is the most effective imaging method?

Sub-question 4: How much more data can be recovered with the optimised forensic collection procedure?

As discussed in Section 3.2.1, the main purpose of this research is to verify SSD forensic capability with current guidelines. Ultimately forensic investigation requires a guideline capable of preserving data from a SSD in a forensically acceptable manner. The outcome of this research will measure current guideline performance against SSD technology tests.

3.2.3 Hypotheses

The hypotheses for this research are as follows.

H1: Current forensic guidelines are incapable of handling SSD.

H2: Data recovered from SSD is not loadable.

H3: Faster imaging method should provide better rate of data recovery.

H4: Combination of certain drives, imaging methods, and write-blockers will facilitate and provide better rate of data recovery.

These hypotheses are derived from the reviewed similar studies. Based on critical literature review in Section 3.2.1, H1 suggests HDD and SSD serve the same purpose as a storage device, but they share nothing in common in terms of architecture and technology. Existing guidelines are developed to maximise the chance of data recovery by deeply understanding the characteristics of a magnetic storage device, and not NAND flash. Therefore it is no sensible to apply methods for magnetic storage devices.

H2 speculates some deleted data can be carved or recovered, but these files are similar to stubs, actual contents are not fully recovered and therefore not loadable. Freeman and Woodward (2009) demonstrated the chance that the recovered files being loadable are zero.

SSD internal processes such as garbage collection and wear-levelling are automated and result in the permanent destruction of evidence. H3 theorise that the longer the acquisition process takes the amount of recoverable data diminishes. Therefore shorter acquisition methods should provide better rates of recovery.

SSD firmware is the core for performance and full of proprietary algorithms. Details of these algorithms are strictly kept confidential, prohibiting forensic investigation from reverse engineering and to reconstruct from data directly extracted from the NAND chips. H4 speculates because firmwares are unique to each SSD model, certain combination of acquisition method should demonstrate effectiveness and therefore provide a better rate of data recovery.

3.2.4 Research Phases

This research consists of four phases, loosely adapting experimental methods from the Association of Digital Forensics, Security and Law (ADFSL) paper discussed in Section 3.1.4 (Bell & Boddington, 2010). Figure 3.2 shows the diagram of the phases.

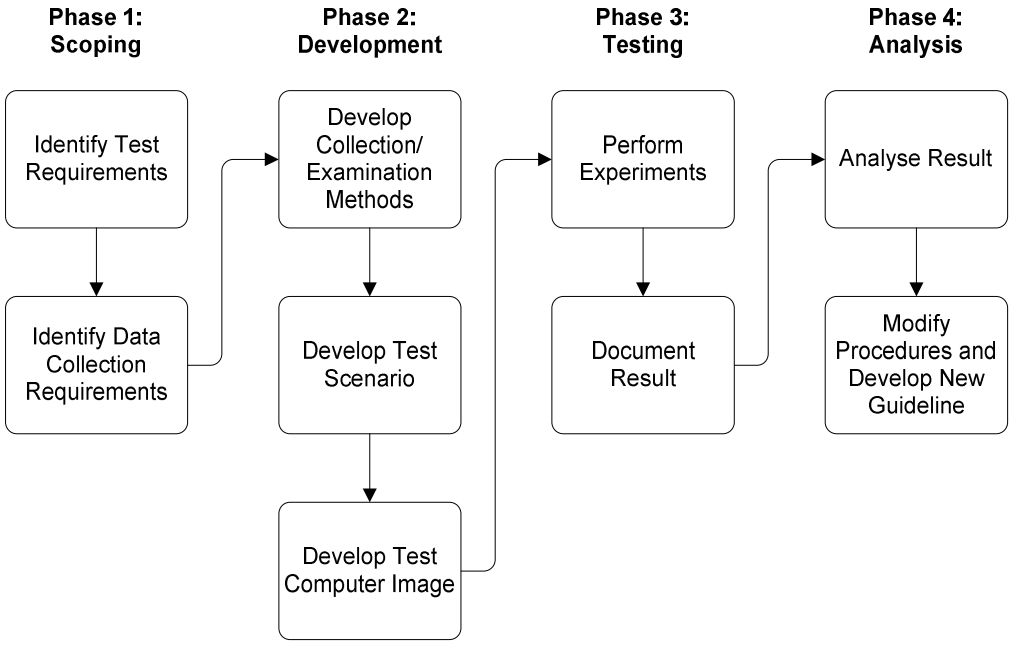

Figure 3.2: Research phases

Phase one is scoping, identifying issues and research gaps from a critical literature review and deriving research questions. Based on the research questions, sample data requirements are suggested to develop the foundations for experiments.

Phase two is a development. Methodologies of collection, acquisition, data recovery tools and techniques are carefully selected and formulated as experimental procedures. A test case scenario and test target computer is prepared for those procedures to be executed, and then the defined data are collected.

Phase three is a testing; carefully planned experiment is tested for data collection and documentation. Changes to the experiment can be made if any issues are found. All activities are logged precisely.

The last phase four is an analysis; which analyse the data and discuss the relevance to the research, then the outcomes are integrated into current forensic guidelines. The combined guideline is composed to maximise the chance of data recovery if a SSD was involved in any computer forensic investigation.

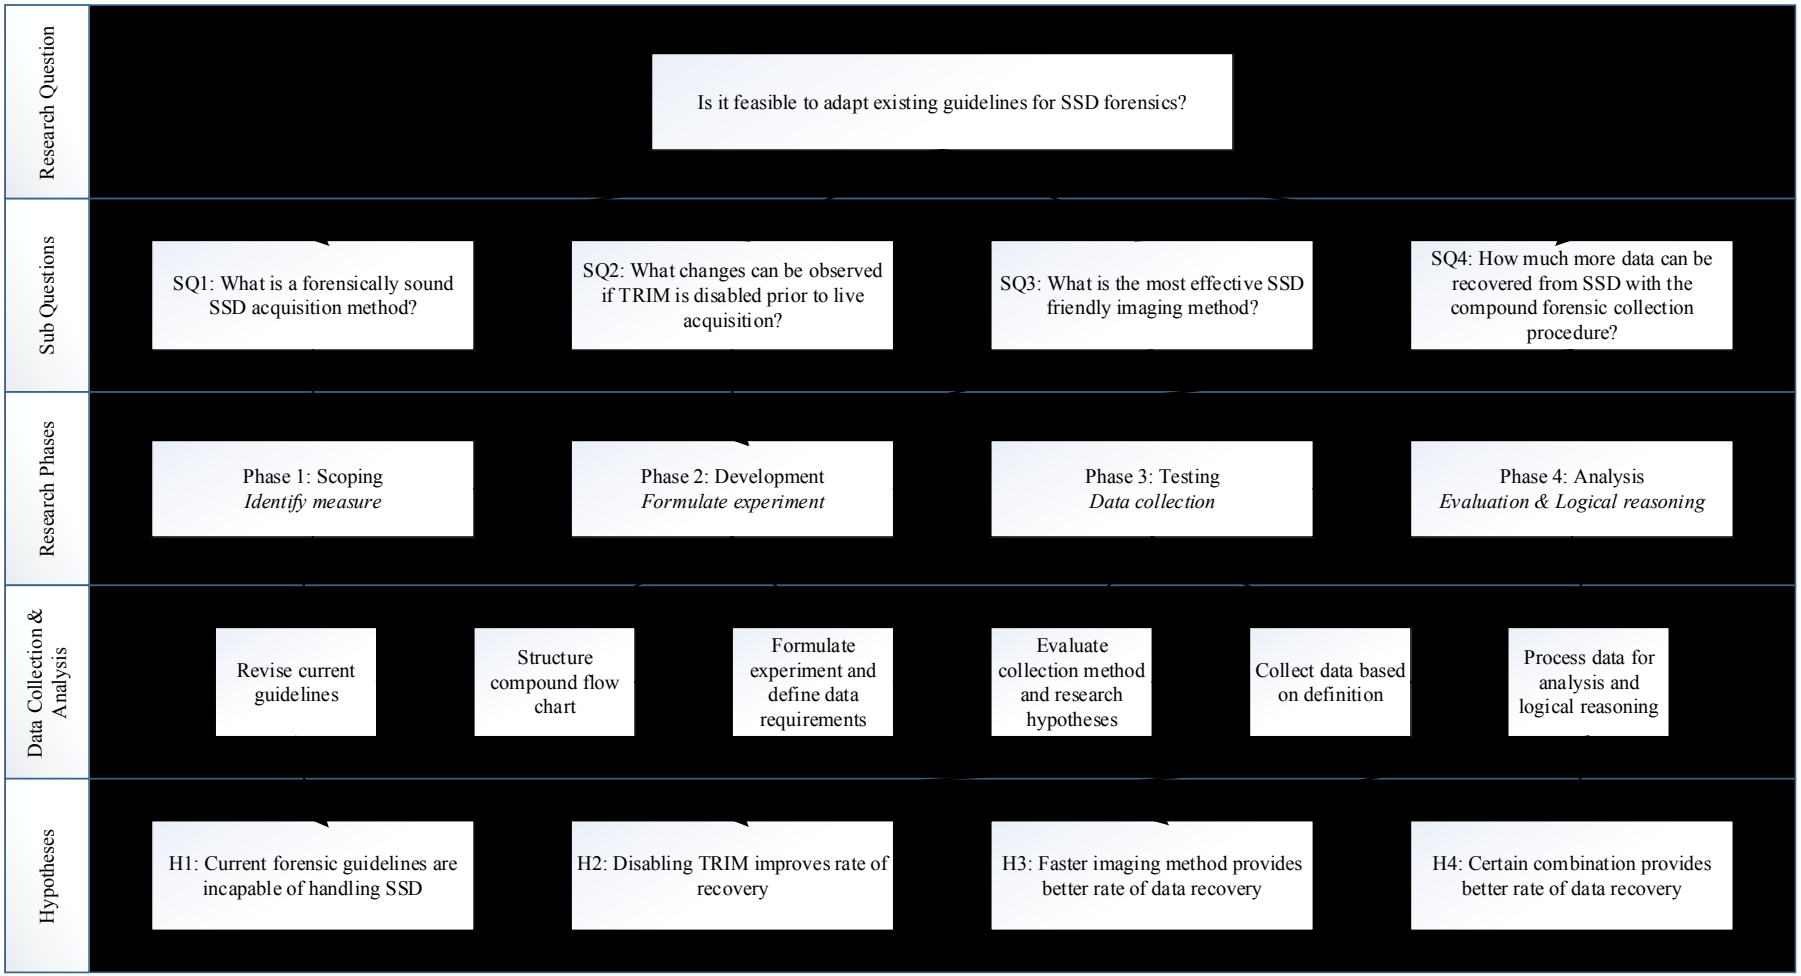

Figure 3.3: Data map

3.2.5 Data Map

The information relationships between the research design components in Section 3.2 are summarised in data map shown in Figure 3.3. The research question and sub-questions are derived from the critical literature review, and research gap identified through review of similar studies. The questions are closely associated with scoping and development phases. These then flow to data collection. The requirements for collected data are carefully formulated, then analysis of the results will provide the outcomes necessary to examine the hypotheses.

3.3 DATA REQUIREMENTS

In the test environment variables are rigorously controlled and the rate of data recovery should purely reflect experimental recovery procedures. The objective of this research is to examine existing guideline's feasibility for forensic data recovery capability on SSD.

Total number of files experimentally deleted and number of data recovered for each test case scenario will provide a base ratio of recovery rates. The specification of deleted files, such as size and file type will enhance data analytic capability. The duration of each experiment is required for the sub-question 3 and H3, which will also create additional dimensions to the research. Testing SSDs are also variables and specifications are required. The data collection method is introduced in Section 3.3.1, and then details of hardware requirements are discussed in Section 3.3.2. Test case scenario will be introduced in Section 3.3.3, and Section 3.3.4 to 3.3.6 will discuss the details of testing methodologies, how the data is processed, and analysed.

3.3.1 Data Collection

The data collection technique is simple. Selected files are deleted then recovered from a computer, and the number of recovered files is counted to calculate the rate of recovery. Environmental factors could significantly affect performance, but deleted files on HDD are generally recoverable. This does not apply to SSD.

This is quantitative research and data is expected to accommodate empirical observation and mathematical analysis. Specific data collection methods are discussed in Section 3.3.3, and processing and analytic methods are discussed in Section 3.3.5 and 3.3.6.

3.3.2 Testing Requirements

Basic requirements for this experiment are listed in Table 3.1.

Table 3.1: List of required items

Item	Purpose
Target computer	Test computer used as target source of storage device
Forensic workstation	Computer used to conduct forensic investigation
HDD	Magnetic storage device used as control variable
SSD	NAND flash storage devices tested for rate of recovery
Write-blockers	Hardware/Software device prevent modifiation to the source
Acquisition tools	Forensic imaging hardware/software
Recovery tools	Recovers deleted data
Test case scenario	Used for multiple test environments and consistency

The target computer is a computer containing data with evidential value and its storage device is acquired for forensic investigation. The Lenovo ThinkPad T430 is selected with Intel Core i5 processor, 4GB of DDR3 random access memory, and running Windows 7 Enterprise 64-bit edition. Various target storage devices are loaded on this computer for testing. In order to retain consistency a Windows image backup will be made and each testing storage devices will have the same image restored. Once the Windows image backup restoration is completed, a planned scenario is followed then acquired for forensic data recovery. This allows the experiment to conduct multiple scenarios on various target drives with the identical data set.

A Forensic workstation is a computer used by examiner for forensic investigations. For these experiments it is a Dell OptiPlex 990 desktop computer with the specifications consisted of Intel Core i7 processor, 4GB of DDR3 random access memory, and running Windows 7 Enterprise 64-bit edition is used.

The HDD is tested on the target computer, for a comparison purpose, and to demonstrate the effectiveness of existing data collection methods on magnetic storage devices. This should also serve as an indicator when the rate of data recovery results are compared. Seagate ST320LT009, 320GB, 7200RPM, SATA 3Gb/s, 16MB cache, 2.5 inch model is selected.

SSDs are tested on the target computer. As discussed in Section 3.1.1, the presence of TRIM creates significant impacts on the ability to recover data. However SSD without TRIM will not be included in this experiment for the following reasons. Firstly all modern SSDs are equipped with TRIM and it is rare

to find a SSD without it. Therefore the experiment presumes all SSDs in the near future will be equipped with TRIM and testing on non-TRIM SSD is not required. Secondly the experiment is not interested in the effect caused by TRIM, but as defined in Section 3.2.2, research sub question one is to identify the effect of disabling TRIM on an operating system prior to an acquisition procedure. Therefore testing on non-TRIM SSD is outside the scope of this experiment. Two recent SSD models are selected. Samsung MZ7PD256HCGM (256GB, SATA 6Gb/s, TLC NAND, TRIM, 512MB DDR2 cache, 2.5 inch) and Intel SSDSC2MH250A2 (250Gb. SATA 6Gb/s, MLC NAND, TRIM, 64MB cache, 2.5 inch).

Write-blockers are used in this experiment for forensic acquisitions. The main purpose of utilising a write-blocker is to prevent any alteration to the original storage device during the acquisition, however as discussed in Section 2.3.3, and 3.1, due to the SSD's internal volition of wiping deleted cells will not be prevented. In order to test the effectiveness of the different acquisition methods, both software and hardware write-blockers shown in Table 3.2 are used in this experiment.

Table 3.2: List of write-blocker used

Manufacture	Model	Version	Type
Tableau	T5		H/W
Guidance Software	Fastbloc (EnCase)	7.07	S/W
SANS	SIFT	2.14	S/W, Bootdisk

Acquisition tools are used to create forensic images of the target storage devices. Depending on the test case scenario, hardware forensic duplicator or software imagers are used for comparison as shown in Table 3.3. Tableau TD2 is a hardware duplicator, device connected to source port is permanently write-blocked, able to create duplicate forensic image of the source device into two destination device at once. It also calculates hash values for verification, and then saves it on the destination drive. Both EnCase and FTK Imager are acquisition software, and EnCase is commercial software that requires a license. FTK Imager is a freeware but does not have a write-block feature. Recovery tools are used to recover deleted files following the testing scenario. Ideally all deleted files are to be recovered, that is possible with HDD if forensic procedures are followed

carefully but unlikely with SSDs. Guidance Software EnCase (version 7.07) will be used to identify deleted files and recovered if possible

Table 3.3: List of forensic acquisition tools

Manufacture	Model	Version	Type	Write-block	Verify
Tableau	TD2		H/W	Yes	MD5, SHA1
Guidance Software	EnCase	7.07	S/W	Yes	MD5, SHA1
Access Data	FTK Imager		S/W	No	MD5, SHA1

3.3.3 Test Case Scenario

Test case scenarios will be used to mimic a real forensic investigation scene as well as properly audit the experiment to minimise involvement of unknown factors. It also supports the experiment to discover answers to the research questions accordingly.

A sample case "possible stolen property" from the NIJ guideline appendix is slightly adjusted and replicated for this experiment. Two main scenarios with six test cases are arranged for each storage device. A summary is shown in Table 3.4.

Test case A and B anticipates the computer power is on when it was discovered. The data recovery performance is tested if live acquisition was implemented. The effectiveness of TRIM during the acquisition is also measured for the sub question 2.

In contrast, test case C to F measures data recovery performance from dead (static) acquisition, where sub-question 3 and 4 are tested by comparing different acquisition methods to seek for better data recoverability.

Table 3.4: Summary of test case scenarios and brief descriptions

Test Case	Scenario	Description
A	Live Acquisition	TRIM disabled then live acquisition
B	Live Acquisition	TRIM enabled then live acquisition
C	Dead Acquisition	H/W write-blocker/duplicator for acquisition
D	Dead Acquisition	Forensic boot disk S/W acquisition
E	Dead Acquisition	S/W write-blocker on forensic workstation and S/W acquisition
F	Dead Acquisition	H/W write-blocker on forensic workstation and S/W acquisition

As the data map in Section 3.2.5 indicated, sub-question 1 and 4 requires comparison of all results in order to determine the effectiveness of collection methods and examine its SSD data recovery capabilities.

3.3.4 Testing Methodology

The experiment is carried out as follows. One HDD and two SSD are tested as a target drive, and each will follow the six test cases. In total 18 forensic images are collected. Each test case starts with sanitising the target drive. Freeman et al. (2009) research on "Testing the efficacy and integrity of secure deletion tools on Solid State Drive" conclude that their experiment shows the DD command is the most effective way of securely wiping SSD. In order to ensure the target drives has no residual data prior to each test, the following command was used in Ubuntu terminal window (Freeman & Woodward, 2009).

dd if=/dev/zero of=/dev/sdb bs=1M

meaning: (if=/dev/zero – overwriting the device with null bytes

of=/dev/sdb – overwriting device is set to the target drive sdb

bs=1M – block size is set to 1 1,048,576 bytes)

The sanitised target drive is installed on the target computer, and Windows recovery CD is used to boot the computer to restore the test computer image.

Power on the restored computer and delete all files under "My Documents" folder, located path of C:\Users\Test_User\Documents, then the test cases are carefully followed for each target drive. Deleted files will be listed and recorded for analysis.

Time taken for each acquisition process is recorded. The same process is repeated until all 18 forensic images are acquired.

Forensic images are copied to the forensic workstation for data recovery. Original images are to be kept securely, and only verified work copies are used for the examination.

A new EnCase file is opened and the acquired image is mounted for data recovery. EnCase features add-ons and specific tools can be added when necessary. Number of recovered files is counted, and a summary is recorded for analysis.

Unallocated and slack spaces are searched, and then copy/un-erase option is used for extraction. Extracted files are also recorded for analysis.

3.3.5 Data Processing

The result from HDD will be a control value in this experiment and the rate of data recovery with SSD is compared to determine its effectiveness. As more of the data can be recovered and closer to the value of control, the effectiveness is better. The rate of recovery is calculated as below.

$$Rate\ of\ Recovery\ (\%) = \frac{Total\ number\ of\ files\ acquired\ +\ Recovered\ files}{\text{Total number of files before deletion}}$$

18 test case results are processed and summarised into a table for analysis. The table contains a matrix of target drives against test cases, and the corresponding rate of recovery as well as acquisition times (duration in minutes) are filled accordingly. Based on the table, a line graph is drawn to graphically represent the results, which should demonstrate correlation between target drives and test cases. A spread sheet is sufficient for this activity.

3.3.6 Data Analysis

The tables and graphs are observed for analysis, which focuses on locating any sign of improvement in data recovery rates amongst the different collection methodologies. The prime variables of interest are any recovery rate value better than another, or closer to the value of control. However other variables such as correlation between rate of recovery and imaging duration and use of different interfaces should also be taken into a consideration.

3.4 LIMITATIONS

There are number of conditions required for this research. Section 3.4 will discuss and specify these limitations. The literature reviews in Section 3.1 suggested lack of widespread understanding is one of current issues with SSD forensics. The research has identified that existing guidelines' capability with SSD is questionable and providing further training may not resolve the root cause for potential failure in forensic processes.

The NAND chip extraction data recovery method discussed in Section 2.3.3 is not considered in this research. This is because there are concerns regarding such extreme modification to hardware that is not a reversible or repeatable process, therefore such a method is considered as not standard forensic practice. The risk of potentially damaging the chip is high, and ability of reproducibility is lost. This is problematic because multiple case scenarios cannot

be tested on the same device. It also requires expert knowledge in data reconstruction as well as electro-circuit modification techniques. In addition, RAID configured SSD is also excluded for similar technical reasons.

Encryption is excluded in this experiment. As discussed in Section 2.3.3, implementation of hardware level encryption mechanisms is getting popular, as well as the number of organisations complying with security standards and enabling data encryption by default is growing rapidly. This indicates ability to break cipher without the original password is essential for computer forensics, and must be developed. But discovering a method for encrypted drives is beyond this thesis's scope and capability and therefore excluded.

The Windows operating system is only used in this research. Although Apple OS X and Linux is growing in popularity the majority, especially in corporate environments, are still using Windows OS. Therefore only Windows 7 operating system image is tested.

Finally, the research scope is limited to data recoverability with SSD in a forensically acceptable manner. Any other forensic guidelines such as cloud computing, smartphone, USB NAND flash drives, and network forensics are not covered and taken into consideration.

3.5 CONCLUSION

Chapter 3 refined the research scope from the foundation developed in the literature review definitions in Chapter 2. Similar studies in relation to SSD data recovery and digital forensic guidelines were critically analysed and research gaps were identified. Research questions were derived from these gaps. The issues identified from the current guidelines in Section 2.4.3 were analysed to form a compound acquisition process flow chart. Hypotheses were drawn for the research questions and research phases were structured. A Data map was used to visually revise the flow of research information and changes were made where necessary.

Data collection methods were adopted based on research discussed in Section 3.1. A small number of test SSD may bias the result, but defined specific data collection and six test case scenarios will allow data analysis from multiple angles and provide sufficient material for the empirical analysis. Chapter 4 will report the research implementation and the results.

Chapter 4

Results

4.0 INTRODUCTION

Chapter 4 reports the results of the testing as specified in Chapter 3. The findings are presented in tables and charts with factual descriptions and empirical analyses where appropriate. In-depth interpretation and evaluation of the results will be discussed in Chapter 5.

Some variations to the methodology defined in Chapter 3 were identified and the rectifications are discussed in Section 4.1. Details of the test environment are reported in Section 4.2, and scenario base test case results are reported in Section 4.3. Finally the results and analyses are presented in Section 4.4. Analyses may include brief interpretations as a linkage to the research questions and hypotheses but the complete discussion will be presented in Chapter 5.

4.1 CHANGES TO SPECIFIED METHODOLOGY

The research method developed in Chapter 3 was tested with a pilot run. The methodology was reassessed and modifications were made. Section 4.1 will discuss identified issues and changes made to the Chapter 3 specification.

4.1.1 Additional SSD Samples

As discussed in Section 3.5, the number of SSD made available was concerning. Larger numbers of sample drives are better, especially knowing that each SSD model could have different firmware, and cell-levels. Similarly different combinations of components have potential to produce different outcomes.

Three additional SSD were obtained but one drive was not suitable for the research methodology. Crucially the M4 128GB SSD was unable to restore from the test Windows 7 recovery image, and only allowing a fresh OS installation. In order to sustain consistency and minimise variables the drive was excluded from the experiment.

The new SSDs are Intel SSDSC2CT120A3 (120GB, SATA 3Gb/s, MLC NAND, TRIM, 256MB cache, 2.5inch) and SanDisk SDSSDHP128G (128GB,

SATA 6Gb/s, MLC NAND, TRIM, DDR2 128MB cache, 2.5inch). Further specifications for tested SSD are listed in Section 4.2.4.

4.1.2 Forensic Boot Disks

Forensic boot disks are typically used for two occasions and a variety of commercial and freeware versions are available. When a storage device is difficult to extract from the computer, the forensic boot disks are used to securely access the storage device without making any changes to it. Secondly the set of forensic software made available on an optical media is an ideal tool set for live acquisition.

An issue identified was that the SANS SIFT boot disk presented in Section 3.3.2 was not capable of creating duplicate image in .E01 format; only .DD the raw image was available. The forensic image consistency became an immediate issue. Although all other tools were capable of creating a raw image format, the use of .E01 was ideal due to its ability to include hash for each image segment.

The number of freely available well known forensic boot disks was critically reviewed and performance tests made for each. The result showed DEFT Forensic Live CD is capable for creating a .E01 image, and therefore was selected as substitute to the SANS SIFT boot disk.

4.1.3 SSD Sanitising Method

Use of a DD command as a SSD sanitising method was originally proposed. Objective of this process is to ensure residual data is purged prior to each test and to eliminate data contamination. The DD command will write zero to all the cells and therefore the SSD will be blank. This process was very slow during the pilot run, taking hours to sterilise a SSD each time.

Tableau has released the latest firmware update revision 7.05 for TD3. TD3 is a unique forensic imager developed and released by Tableau in 2011. The latest firmware update included the ability to “Secure Erase” a SSD. The TD3 Secure Erase was tested on the every test SSD, and the performance was acceptable. Each SSD took around ten seconds to complete the wipe. Each wiped SSD was mounted on EnCase and the forensic software was unable to locate any residual data.

As discussed in Section 2.3.2 and 3.1.3, secure erase, or formally known as the ATA Secure Erace command, is not physically over writing or purging

electrons from the cells and making the device blank. It destroys old encryption keys and replaces them with another. This makes existing data invalid and a collection of scrambled 0s and 1s. Besides, as far as a forensic tool is concerned, EnCase was unable to detect such theoretical residual and therefore it was verified that use of the Secure Erace is forensically acceptable method of sanitising SSD.

4.2 TESTING ENVIRONMENT

The tests were run in the Auckland University of Technology forensic lab. Section 4.2 describes the testing environment including available computers, storage devices, write blockers, and digital forensic software used to recover deleted data.

4.2.1 Test Computer

One laptop computer was used as a target computer; five testing storage devices are used on the computer. The target computer specifications are listed below.

Table 4.1: Test Computer Specification

Manufacturer	Dell Inc.
Product Make	Latitude E6410
Motherboard	Dell Inc. 0667CC
BIOS Info	AT/AT COMPATIBLE | 08/10/10 | DELL - 6222004
SM BIOS	A05
Memory (RAM)	3894 MB
CPU Info	Intel(R) Core(TM) i5 CPU M 560 @ 2.67GHz
CPU Speed	2657.4 MHz
Network Adapters	Bluetooth Device (Personal Area Network) #5 | Intel(R) Centrino(R) Advanced-N 6200 AGN | Intel(R) 82577LM Gigabit Network Connection
CD / DVD Drives	HL-DT-STDVD-ROM DU30N
Time Zone	New Zealand Standard Time

4.2.2 Forensic Computer

One desktop computer was used as a forensic computer. The following table lists the specifications.

Table 4.2: Forensic Computer Specification

Manufacturer	Dell Inc.
Product Make	OptiPlex 990
Motherboard	Dell Inc. 0VNP2H
SM BIOS	A17

Windows	Windows 7 Enterprise Edition (64-bit) SP1 (Build 7601)
Internet Explorer	9.10.9200.16686
Memory (RAM)	16266 MB
CPU Info	Intel(R) Core(TM) i7-2600 CPU @ 3.40GHz
CPU Speed	3440.5 MHz
Display Adapters	Intel(R) HD Graphics | Intel(R) HD Graphics
Network Adapters	Intel(R) 82579LM Gigabit Network Connection
CD / DVD Drives	TSSTcorpDVD+-RW TS-H653G
USB Controllers	2 host controllers.
Firewire (1394)	Not Detected
BIOS Info	AT/AT COMPATIBLE | 03/14/13 | DELL - 6222004
Time Zone	New Zealand Standard Time

4.2.3 Test System Image

As defined in Section 3.3.1, the purpose of using system image is to minimise variables. Windows image restoration allows the target computer with different storage devices to be logically the same condition. Table 4.3 presents the specification of the Windows image used.

Table 4.3: Details of test operating system configurations

Windows Version	Windows 7 (Build 7601)
Edition	Enterprise Edition (64-bit)
Internal Version	6.1
Bit Size	64 Bit
Service Packs	Service Pack 1
Internet Explorer	9.10.9200.16721
Windows Language	English (New Zealand)
Windows Folder	C:\Windows\
Boot Mode	Normal Boot

4.2.4 Storage devices

Table 4.4 to 4.8 list the details of the storage devices tested.

Table 4.4: Test Drive "HDD-01" specification

Model	HTS541680J9SA00
Manufacture	Hitachi
Firmware	SB2IC7JP
Serial	SB2241SGFR5WUE
Bus Type	S-ATA/150 (over USB 3.0)
ATA Standard	ATA/ATAPI-7 T13 1532D version 1
Transfer Mode	S-ATA 150
Size	80.0 GB (74.5 GiB) (48bit-LBA)

Buffer	7.7 MB
Rotation Rate	5,400 RPM
Features	48bit LBA / NCQ / HPA / DCO / GPL
Transfer Feature	S-ATA I , UDMA 100 / 66 / 44 / 33 / 25 / 16 , PIO Mode 4 / 3
Security Feature	supported / Master Password Capability: High
AAM	Yes (Current Level: 0xFE) / Enabled
APM	Yes (Current Level: 0x80) / Enabled

Table 4.5: Test Drive "SSD-01" specification

Model	SS INTEL SSDSC2MH250A2
Manufacture	Intel
Firmware	PWG4
Serial	LNEL116200U1250DGN
Bus Type	S-ATA/600 (over USB 3.0)
ATA Standard	ATA8-ACS revision 2d
Transfer Mode	S-ATA 600 (S-ATA Rev 3.0)
Size	250.1 GB (232.9 GiB) (48bit-LBA)
Buffer	
Rotation Rate	SSD (Solid State Device)
Features	48bit LBA / NCQ / HPA / DCO / TRIM / GPL
Transfer Feature	S-ATA III / II / I , UDMA 133 / 100 / 66 / 44 / 33 / 25 / 16 , PIO Mode 4 / 3
Security Feature	supported / enhanced security erase / Master Password Capability: High
AAM	No
APM	No

Table 4.6: Test Drive "SSD-02" specification

Model	SSD 840 PRO Series
Manufacture	Samsung
Firmware	DXM04B0Q
Serial	S12RNEAD127172F
Bus Type	S-ATA/600 (over USB 3.0)
ATA Standard	ACS-2 revision 4c
Transfer Mode	S-ATA 600 (S-ATA Rev 3.1)
Size	256.1 GB (238.5 GiB) (48bit-LBA)
Buffer	
Rotation Rate	SSD (Solid State Device)
Features	48bit LBA / NCQ / HPA / DCO / TRIM / GPL
Transfer Feature	S-ATA III / II / I , UDMA 133 / 100 / 66 / 44 / 33 / 25 / 16 , PIO Mode 4 / 3
Security Feature	supported / enhanced security erase / Master Password Capability: High
AAM	No
APM	No

Table 4.7: Test Drive "SSD-03" specification

Model	SS INTEL SSDSC2CT120A3
Manufacture	Intel
Firmware	300i
Serial	CVMP2163016T120BGN
Bus Type	S-ATA/300 (over USB 3.0)
ATA Standard	ACS-2 revision 3
Transfer Mode	S-ATA 300 (S-ATA Rev 3.0)
Size	120.0 GB (111.8 GiB) (48bit-LBA)
Buffer	
Rotation Rate	SSD (Solid State Device)
Features	48bit LBA / NCQ / HPA / TRIM / GPL
Transfer Feature	S-ATA II / I , UDMA 133 / 100 / 66 / 44 / 33 / 25 / 16 , PIO Mode 4 / 3
Security Feature	supported / enhanced security erase / Master Password Capability: High
AAM	No
APM	APM: Yes (Current Level: 0xFE) / Enabled

Table 4.8: Test Drive "SSD-04" specification

Model	SDSSDHP128G
Manufacture	SanDisk
Firmware	X211200
Serial	1.3125E+11
Bus Type	S-ATA/600 (over USB 3.0)
ATA Standard	ATA8-ACS revision 6
Transfer Mode	S-ATA 600 (S-ATA Rev 3.0)
Size	128.0 GB (119.2 GiB) (48bit-LBA)
Buffer	
Rotation Rate	SSD (Solid State Device)
Features	48bit LBA / NCQ / HPA / DCO / TRIM / GPL
Transfer Feature	S-ATA III / II / I , UDMA 133 / 100 / 66 / 44 / 33 / 25 / 16 , PIO Mode 4 / 3
Security Feature	supported / enhanced security erase / Master Password Capability: High
AAM	No
APM	APM: Yes (Current Level: 0x80) / Enabled

4.2.5 Write-blockers

The following table lists the write-blockers used (Table 4.9).

Table 4.9: Write-blocker specifications

Manufactu	Model	Firmware	Serial	Type

re				
Tableau	Forensic Duplicator TD2	v7.03	01D220AA 1206	Hardware
Tableau	Forensic SATA Bridge T3u	v5.20	V006C026590	Hardware
Guidance Software	FastBloc SE	-	-	Software

4.2.6 Forensic Software

The following forensic software was used to create the forensic images (Table 4.10).

Table 4.10: Details of forensic software

Manufacture	Model	Version
Guidance Software	EnCase® Forensic	6.19.6
DEFT Association	DEFT Forensic Live CD	6.1
AccessData	FTK Imager Lite	3.1.1

4.3 DATA COLLECTION

As discussed in section 3.3.1, the main purpose of the data collection is to examine the best acquisition method for preserving deleted data and to maximise the chance of data recovery when forensic image is created with SSD. The evidence destruction behaviour triggered by a SSD built-in optimisation process was discussed in Section 2.3.3. This revised data collection methods are designed to collect the sample rate of data recovery from four test SSDs. Six tests designed in Section 3.2 are used for each SSD, and the tests will follow the methodology discussed in section 3.3.4. Details of these methodologies and test case scenarios are discussed in Section 4.3.1 to 4.3.6.

4.3.1 Case A

The storage device (target drive) was wiped, and then the test system image was restored using Windows image recovery. Once the image restoration was completed, all the folders and files stored under “My Document” folder were deleted. After five minutes of waiting time, a FTK imager CD was loaded, and an external USB storage device was connected. The “Physical” target drive was selected as an image source, and then the destination was set to the external device. The following image file naming convention was used.

SSD-XX-A.e01 (XX = Testing SSD numbers, e.g. SSD-01-A)

The beginning and completion of acquisition times were recorded to calculate the duration. Verified images were copied to the forensic workstation for forensic analysis. Deleted files were analysed in particular.

EnCase was used to recover files from the forensic images. The Images were mounted using EnCase, target partition was selected for the "Recover Folders" option. The number of items recovered from the "My Document" folder is recorded. Recovered items are sorted by the following file types.

Table 4.11: File extension and types used for recovery test

PDF	Portable Document Format
FLV	Flash Video Format
GIF	Graphics Interchange Format
HTML	Hyper Text Markup Language Format
JPEG	Joint Photographic Expert Group Format
DOCX	Microsoft Office Open XML Format

The number of recovered items for each file type was recorded, and then opened to test its complete recovery. The number of files being able to open was recorded.

4.3.2 Case B

The storage device (target drive) was wiped, and then the test system image was restored using Windows image recovery. Once the image restoration was completed, all the folders and files stored under "My Document" folder were deleted. After five minutes of waiting time, TRIM was disabled by running the following command line from command prompt.

fsutil behavior set DisableDeleteNotify 1

FTK imager CD was loaded, and an external USB storage device was connected. The "Physical" target drive was selected as an image source, and then destination was set to the external device. The following image file naming convention was used.

SSD-XX-B.e01 (XX = Testing SSD numbers, e.g. SSD-01-B)

The same procedure from Case A from this point onwards was followed. The duration, verification, recovery, and recovery verification were tested.

4.3.3 Case C

Storage device (target drive) was wiped, and then the test system image was restored using Windows image recovery. Once the image restoration completed, all the folders and files stored under “My Document” folder were deleted.

After five minutes of waiting time, the laptop was force hardware shutdown by holding power button for seven seconds. Power and battery was removed, and then the target drive was removed from the laptop. Removed target drive was acquired with Tableau Forensic Duplicator TD2, by following the steps below. Note that in this research, the term “source drive” refers to the origin of data acquired, and “target drive” refers to the destination of data copied to.

1. Membrane Power Switch. Confirm TD2 is OFF.
2. Connect the target drive to "Source Drive Interface" with SATA cable to TD2.
3. Destination drive was connected to "Destination Drive Interface".
4. Connect TP4 power supply to TD2 DC In.
5. Power on the TD2, source drive, and destination drive by pressing the TD2 membrane power switch.
6. From the Main menu, use the arrow keys to navigate to Duplicate Disk > Disk-to-File (Menu 1.2).
7. Select EnCase Format (.E01)
8. Enter Case ID as "C"
9. Enter Case Note as *SSD-XX-C* (XX = Test Storage ID, e.g. SSD-01-C)
10. Enter File Name as *SSD-XX-C.e01*
11. Imaging begins and progress reports appear on the LCD.
12. When the process is completed, check the log to ensure there was no issues during imaging.

The same procedure from Case A from this point onwards was followed. The duration, verification, recovery, and recovery verification were tested.

4.3.4 Case D

Storage device (target drive) was wiped, and then the test system image was restored using Windows image recovery. Once the image restoration completed, all the folders and files stored under "My Document" folder were deleted. After five minutes of waiting time, the laptop was force hardware shutdown by holding power button for seven seconds. Power and battery was removed, and then DEFT forensic live CD is placed in the optical drive. The external destination drive was connected to the laptop USB port.

Then plug the power back into the laptop, and press F12 on the keyboard to select boot device. The CD/DVD drive was selected and DEFT operating system appeared. Select "live" and then typed "gui-deft" to launch graphic user interface. Navigate the interface and select "guymager" from disk imaging menu.

The target drive was selected from the list of storage devices, right click and select acquire. The details were filled as presented in Table 4.12.

Table 4.12: Example of details filled for test case D acquisition process

Criteria	Details
Split Size (MiB)	*2047*
Case Number	*D*
Evidence Number	*SSD-XX-D*
Image Directory	*External Target Storage Device*
Image Filename	*SSD-XX-D.e01*
Calculate MD5	*Enable*
Verify Image	*Enable*

The same procedure from Case A from this point onwards was followed. The duration, verification, recovery, and recovery verification was tested.

4.3.5 Case E

Storage device (target drive) was wiped, and then the test system image was restored using Windows image recovery. Once the image restoration was completed, all the folders and files were stored under the "My Document" folder and deleted. After five minutes of waiting time, the laptop was forced into hardware shutdown by holding power button for seven seconds. The power and battery was removed, and then the target drive was removed from the laptop. The removed target drive was acquired with Guidance Software EnCase, by following the steps below.

1. On the forensic workstation, opened EnCase version 6.19.6, and created a new case "E".
2. Navigated to Tools, FastBloc SE, and ensured "write Blocked" mode was selected.
3. Attached the target drive to the workstation USB port. The target drive appeared in the list of volumes and confirmed the drive write block status was showing "Yes".
4. Add the target drive as a "Local Device", select the write-blocked target device and click finish.
5. Right click the target drive (physical) in the table pane, select Acquire, and fill the options as shown in Table 4.13.

Table 4.13: Example of details filled for test case E acquisition process

Criteria	Details	Example
Name:	*SSD-XX-E*	SSD-01-E
Evidence Number	*SSD-XX-E*	SSD-01-E
File Segment Size	*2000*	
Output Path	*External Target Storage Device*	

The same procedure from Case A from this point onwards was followed. The duration, verification, recovery, and recovery verification test.

4.3.6 Case F

Storage device (target drive) was wiped, and then the test system image was restored using Windows image recovery. Once the image restoration was completed, all the folders and files stored under the “My Document” folder were deleted.

After five minutes of waiting time, the laptop was force hardware shutdown by holding power button for seven seconds. The power and battery was removed, and then the target drive was removed from the laptop. The removed target drive was acquired with Guidance Software EnCase and Tableau T3u write-blocker, by following the steps below.

1. On the forensic workstation, opened EnCase version 6.19.6, and created a new case "F".

2. Navigated to Tools, FastBloc SE, and ensured "write Blocked" mode was selected.
3. Attached the target drive to Tableau T3u write-blocker by following steps below, and then connected to the workstation using USB cable.
    a. Confirmed T3u power switch is in the off position.
    b. Connect T3u to the SATA subject (target) drive.
    c. Connect subject SATA drive to T3u DC power out.
    d. Connect host (the forensic workstation) to T3u using USB 2.0 cable.
4. Add the target drive as a "Local Device", select the write-blocked target device and click finish.
5. Right click the target drive (physical) in the table pane, select Acquire, and fill the options as shown in Table 4.14.

Table 4.14: Example of details filled for test case F acquisition process

Criteria	Details	Example
Name:	*SSD-XX-F*	SSD-01-F
Evidence Number	*SSD-XX-F*	SSD-01-F
File Segment Size	*2000*	
Output Path	*External Target Storage Device*	

The same procedure from Case A from this point onwards was followed. The duration, verification, recovery, and recovery verification were tested.

4.4 TEST RESULTS

Results of the data collection and analysis are provided in Section 4.4. Section 4.4.1 will group the results by the case scenarios, demonstrate which SSD has exceled in each test case and correlation amongst file types, duration, and successful rate of recovery will be reported. Successful recovery means that files are recovered as well as they are loadable so that the contents can be viewed.

Section 4.4.2 will group the same results by SSDs, identifying average rates of recovery for each SSDs and identify outstanding exceptions. This analysis will assist identify if certain drives have vigorous background wiping which prohibits any attempt of recovery.

The overall rate of recovery by file types and durations will also be analysed in Section 4.4.3. This will demonstrate overall data recovery capabilities and behavioural characteristics when the current guidelines were applied to SSDs.

4.4.1 Test Results Analysis Grouped by Test Drives

Section 4.4.1 analyses collected data grouped by the test drives. Figure 4.1 to 4.3 illustrated the combined results in a clustered chart. The rate of recovery was compared across the test drives in Figure 4.1. The average acquisition time taken for each test drives are shown in Figure 4.2 and 4.3.

Table 4.15 to 4.19 analysed the test results grouped by the test drives. The column labelled "Test" in the tables describes the test case identifier. The column with file extensions (".pdf", ".flv", ".gif", ".html", ".jpg", and ".docx") provides the number of files recovered fully or partially readable. The column "Rate" provides percentage of recovery out of the deleted items. The column "Time" provides the duration required to perform a forensic image acquisition. Note that the test case C only produced time to the nearest minute.

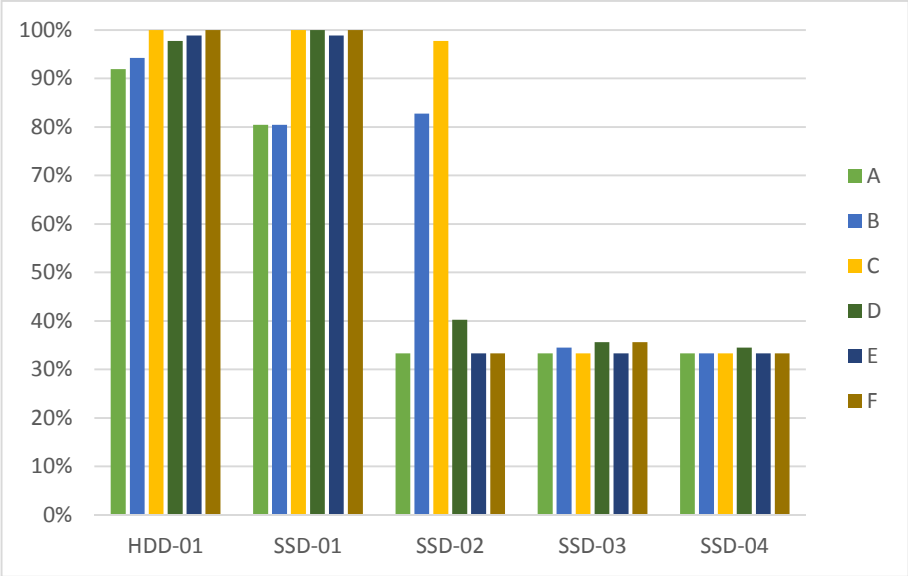

Figure 4.1: Test drives rate of recovery are displayed in clustered column graph

SSD-01 solely outperformed amongst others, the successful rate of recovery is almost equivalent to HDD-01. While the rates for SSD-03 and 04 were just above 30% average, test case B and C enabled SSD-02's ability to recover the deleted data was more than double.

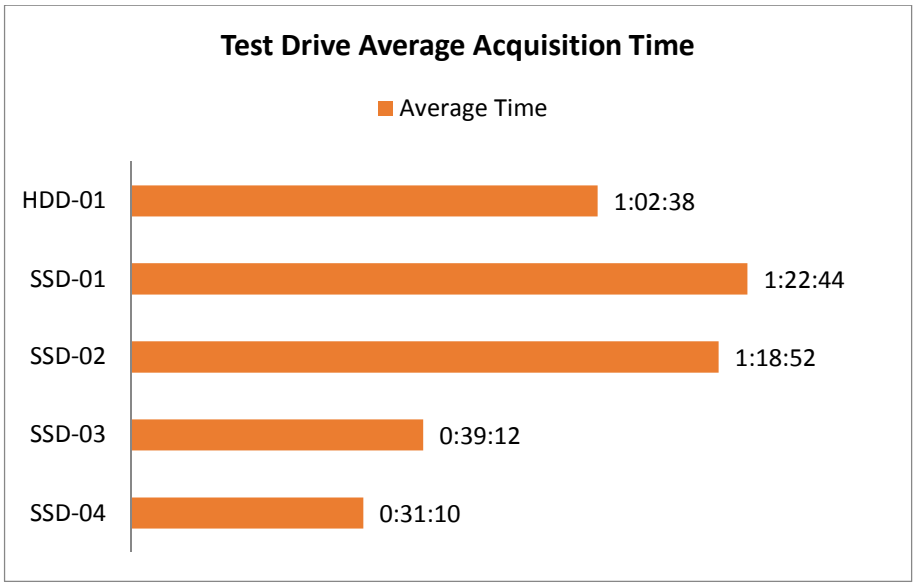

Figure 4.2: Test drive average acquisition times are displayed in clustered bar graph

Figure 4.2 shows the average acquisition time required for each drive. However, although the chart is accurately plotting the average times, comparison between the drives are deceptive and require a data normalisation. The storage size for each test drives are different, and theoretically a SSD with a larger storage capacity should take a longer time because of the difference in volume size. Therefore to compare the acquisition data, the acquisition time is divided by the capacity size in nearest GB.

$$Normalised\ Acquisition\ Time = \frac{Actual\ Acquisition\ Time}{Device\ Storage\ Capacity\ (GB)}$$

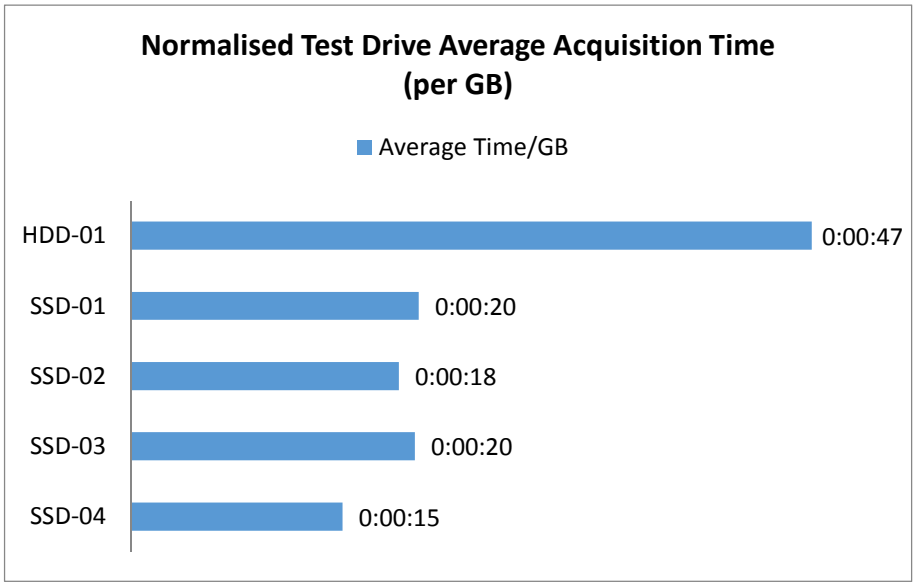

Figure 4.3: Normalised average acquisition times for test drives in clustered bar graph

Figure 4.3 shows the visually comparable normalised average acquisition time for the test drives. The fastest average acquisition time was recorded on SSD-04, showing 25% faster speed when compared with SSD-01 and 03.

Table 4.15: Test results for HDD-01 (Value of Control)

Test	.pdf	.flv	.gif		.html	.jpg	.docx	Rate	Time
A	7	2	33		4	33	1	92%	0:47:50
B	9	1	33		4	33	2	94%	0:48:03
C	9	4	33		4	35	2	100%	0:40:00
D	9	4	33		4	33	2	98%	0:37:40
E	9	4	33		4	34	2	99%	2:30:50
F	9	4	33		4	35	2	100%	0:51:25

Value of control collected from HDD-01 were generally as expected and satisfactory. As Table 4.15 indicates, live acquisition provided slightly less rate of recovery than dead acquisition. TRIM command does not affect non-SSD drive therefore the same rate was expected between test A and B, however HDD-01 demonstrated slightly better rate of recovery when TRIM was disabled.

Table 4.16: Test results for SSD-01

Test	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time
A	0	1	32	4	33	0	80%	0:46:14
B	0	2	32	4	32	0	80%	0:46:20
C	9	4	33	4	35	2	100%	0:33:00
D	9	4	33	4	35	2	100%	0:40:10
E	9	4	33	4	34	2	99%	2:40:26
F	9	4	33	4	35	2	100%	2:50:12

The overall rate of recovery from SSD-01 was unexpected, well exceeding expectations and outperformed others. Live acquisition on SSD-01 was less effective, unable to recover none of .docx and .pdf files and .flv was less successful as well.

Table 4.17: Test results for SSD-02

Test	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time
A	0	0	27	2	0	0	33%	0:42:33
B	0	0	31	4	35	2	83%	0:42:27
C	9	4	31	4	35	2	98%	0:33:00
C-1	9	4	32	4	34	2	98%	0:17:00
D	0	0	33	2	0	0	40%	0:40:42
E	0	0	27	2	0	0	33%	2:47:25
F	0	0	27	2	0	0	33%	2:27:05

Six original test case results and one additional test results are shown in Table 4.20. The effect of disabling TRIM appeared in the result between test A and B, where a 50% increase in rate of recovery was monitored. However, dead acquisition tests were only recovering .gif and .html files with the exception of test C that recovered 98% of deleted files. Potential human error was suspected and the same test was performed in C-1.

The same result was obtained but the acquisition time was almost halved. In summary, SSD-02 demonstrated fastest acquisition speed in test C, and the highest recovery rate of produced in test C, while a longer acquisition time resulted in lower recovery rates. This result exhibited faster acquisition speeds and has direct correlation to better rate of recovery.

Table 4.18: Test results for SSD-03

Test	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time
A	0	0	27	2	0	0	33%	0:27:22
B	0	0	27	3	0	0	34%	0:27:19
C	0	0	27	2	0	0	33%	0:15:00
D	0	2	27	2	0	0	36%	0:20:34
E	0	0	27	2	0	0	33%	1:09:00
F	0	2	27	2	0	0	36%	1:16:00

Table 4.19: Test results for SSD-04

Test	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time
A	0	0	27	2	0	0	33%	0:28:25
B	0	0	27	2	0	0	33%	0:27:59
C	0	0	27	2	0	0	33%	0:17:00
D	0	0	27	3	0	0	34%	0:21:35
E	0	0	27	2	0	0	33%	0:13:00
F	0	0	27	2	0	0	33%	1:19:00

The results from SSD-03 and 04 were similar where none of .pdf, .jpg, and .docx were recovered at all. 27 .gif images and 2 .html files were loadable therefore the results were better than the research hypothesis expected, but promoted new question of why only certain files were always recovered while some file types were harder to be recovered.

4.4.2 Test Results Analysis Grouped by Test Case

Figure 4.4 to 4.6 illustrates the combined results in clustered charts. Clustered charts allows value comparison across categories, in this case test cases against SSDs are compared. The rate of recovery is compared across the test cases in

Figure 4.4. The average acquisition time taken for each test case is shown in Figure 4.5. A chart for the normalised average time is also available in Figure 4.6.

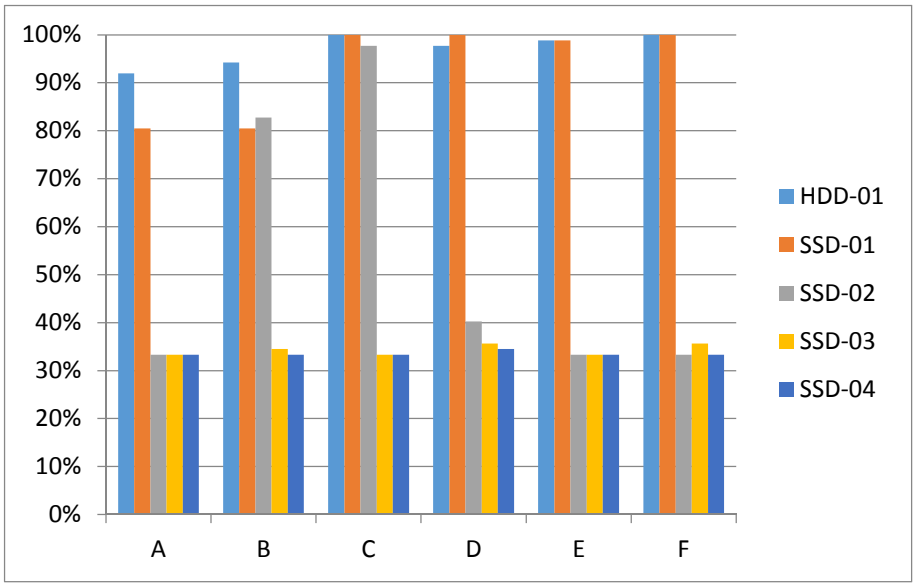

Figure 4.4: Results displayed in clustered column graph grouped by test case

Table 4.20 to 4.25 summarise the test results grouped by the test cases. The column labelled "Drive" in the tables describes the test drive identifier. The column with file extensions (".pdf", ".flv", ".gif", ".html", ".jpg", and ".docx") gives the number of files recovered fully or partially readable. The column "Rate" gives percentage of recovery out of the deleted items. The column "Time" gives the duration required to perform forensic image acquisition. Lastly the column "Normalised" shows normalised acquisition time. Note that the test case C was only able to measure duration to the nearest minutes.

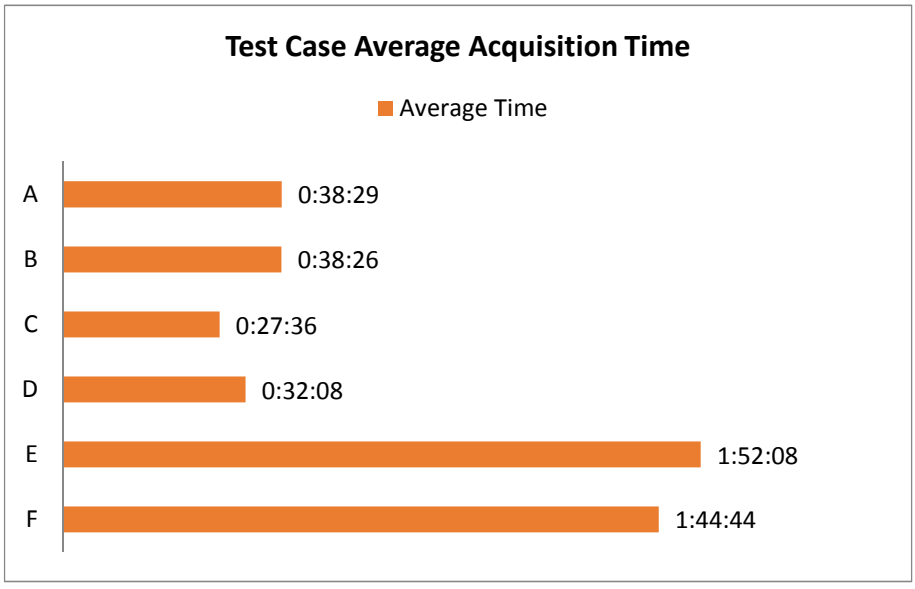

Figure 4.5: Results displayed in clustered bar graph grouped by test cases

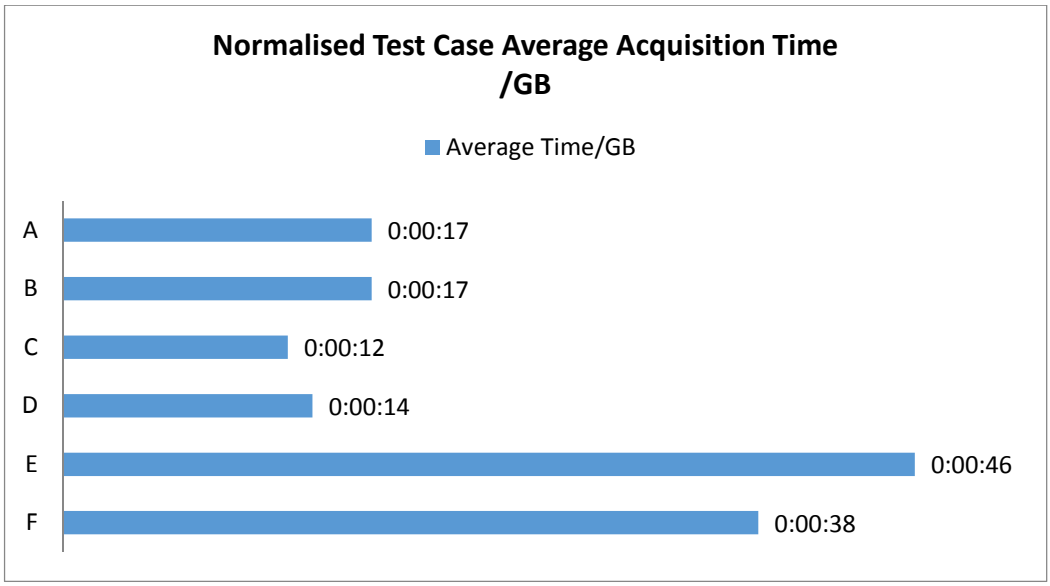

Figure 4.6: Normalised results displayed in clustered bar graph grouped by test cases

In test case A, SSD-01 managed to recover 80% of the deleted files while the average recovery rate for case A was 45%. The significant difference was due to the success of 33 out of 35 .jpg image file recovery. Acquisition time for SSD-03 was the longest in the group which took over 14 seconds per GB, and the shortest was SSD-02's 10 seconds per GB. However the rate of recovery was both 33%. Without normalisation, SSD-01 took the longest over 46 minutes acquisition time, but recovered 80% of the deleted data.

Table 4.20: Test case "A" results

Drive	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time	Normalised
HDD-01	*7*	*2*	*33*	*4*	*33*	*1*	*92%*	*0:47:50*	0:00:36
SSD-01	0	1	32	4	33	0	80%	0:46:14	0:00:11
SSD-02	0	0	27	2	0	0	33%	0:42:33	0:00:10
SSD-03	0	0	27	2	0	0	33%	0:27:22	0:00:14
SSD-04	0	0	27	2	0	0	33%	0:28:25	0:00:13

SSD-02 showed better rate of recovery than SSD-01 in test case B, while SSD-03 and SSD-04 were showing no sign of improvement when compared with test A. Only SSD-01 was able to recover two .flv video files, and only SSD-02 was able to recover two .docx document files. Both SSD-01 and SSD-02 acquisition time were longer than overall average of 36 minutes. Normalised acquisition time remain similar, ranging from 10 to 14 seconds per GB for all the drives.

Table 4.21: Test case "B" results

Drive	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time	Normalised
HDD-01	*9*	*1*	*33*	*4*	*33*	*2*	*94%*	*0:48:03*	0:00:36
SSD-01	0	2	32	4	32	0	80%	0:46:20	0:00:11
SSD-02	0	0	31	4	35	2	83%	0:42:27	0:00:10
SSD-03	0	0	27	3	0	0	34%	0:27:19	0:00:14
SSD-04	0	0	27	2	0	0	33%	0:27:59	0:00:13

SSD-01 and SSD-02 continued to perform well in test case C, SSD-01 recovered all the deleted files. Acquisition time was shorter than the previous live acquisition test cases A and B, the normalised average time taken was 8 seconds per GB. SSD-03 and 04 recovery rates remained unchanged of below 34%.

Table 4.22: Test case "C" results

Drive	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time	Normalised
HDD-01	*9*	*4*	*33*	*4*	*35*	*2*	*100%*	*0:40:00*	0:00:30
SSD-01	9	4	33	4	35	2	100%	0:33:00	0:00:08
SSD-02	9	4	31	4	35	2	98%	0:33:00	0:00:08
SSD-03	0	0	27	2	0	0	33%	0:15:00	0:00:07
SSD-04	0	0	27	2	0	0	33%	0:17:00	0:00:08

SSD-02 struggled in test case D, while SSD-01 continued to recover 100% even though HDD-01 (value of control) recovered 2 out of 35 .jpg image files less than 98%. Overall normalised average duration was just over 10 seconds per GB, a 20% increase from test case C but average rate of recovery fell down to 53% while test case C average was 66%.

Table 4.23: Test case "D" results

Drive	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time	Normalised
HDD-01	*9*	*4*	*33*	*4*	*33*	*2*	*98%*	*0:37:40*	0:00:28
SSD-01	9	4	33	4	35	2	100%	0:40:10	0:00:10
SSD-02	0	0	33	2	0	0	40%	0:40:42	0:00:10
SSD-03	0	2	27	2	0	0	36%	0:20:34	0:00:10
SSD-04	0	0	27	3	0	0	34%	0:21:35	0:00:10

Test case E was the slowest in terms of acquisition time, the normalised average was 30 seconds per GB, with the exception of SSD-04 which outperformed and completed 5 times faster than the average speed. The recovery rates were not impressive, with an average of 50%. Only SSD-01 exceled in the test. Other SSDs recovered 33%.

Table 4.24: Test case "E" results

Drive	.pdf	.flv	.gif	.html	.jpg	.pdf	Rate	Time	Normalised
HDD-01	*9*	*4*	*33*	*4*	*34*	*2*	*99%*	*2:30:50*	0:01:53
SSD-01	9	4	33	4	34	2	99%	2:40:26	0:00:39
SSD-02	0	0	27	2	0	0	33%	2:47:25	0:00:39
SSD-03	0	0	27	2	0	0	33%	1:09:00	0:00:35
SSD-04	0	0	27	2	0	0	33%	0:13:00	0:00:06

The results in test case F were similar to test case E. SSD-01 spent over 170 minutes for acquisition and recovered the all deleted file. In contrast, SSD-02's acquisition time was just over 147 minutes but only recovered 33%. SSD-03 and 04 continued to struggle with their recovery rate, although the durations were almost half compared with SSD-01 and 02. This contradicts with the result analysed in Section 4.4.1, where supporting evidence of faster acquisition leads to a better recovery rate as examined from the SSD-02 results.

Table 4.25: Test case "F" results

Drive	.pdf	.flv	.gif	.html	.jpg	.docx	Rate	Time	Normalised
HDD-01	*9*	*4*	*33*	*4*	*35*	*2*	*100%*	*0:51:25*	0:00:39
SSD-01	9	4	33	4	35	2	100%	2:50:12	0:00:41
SSD-02	0	0	27	2	0	0	33%	2:27:05	0:00:34
SSD-03	0	2	27	2	0	0	36%	1:16:00	0:00:38
SSD-04	0	0	27	2	0	0	33%	1:19:00	0:00:37

4.4.3 Test Results Analysis in Specific Purpose Charts

In addition to section 4.4.1 and 4.4.2, Figure 4.7 to 4.9 show relationships of individual items to the whole, illustrating the contribution of each result to a total across the test drives. Figure 4.7 shows the rate of recovery for each file types across the test drives in stacked column chart to demonstrate the relationships, while Figure 4.8 and Table 4.26 illustrates the statistics of the average file size and associated recovery rate for each file type. Figure 4.9 shows a correlation between the total rate of recovery against the acquisition time for the test drives.

As identified in analyses Section 4.4.1 and 4.4.2, the recovery rates in Figure 4.7 shows a trend that .pdf, .jpg and .docx were harder to recover while some .gif and .html were consistently recovered from all test drives. Interestingly .flv video files were still recoverable in some cases. As far as file types were concerned, in-depth analysis on why some files can be consistently recovered while others failed needs to be assessed.

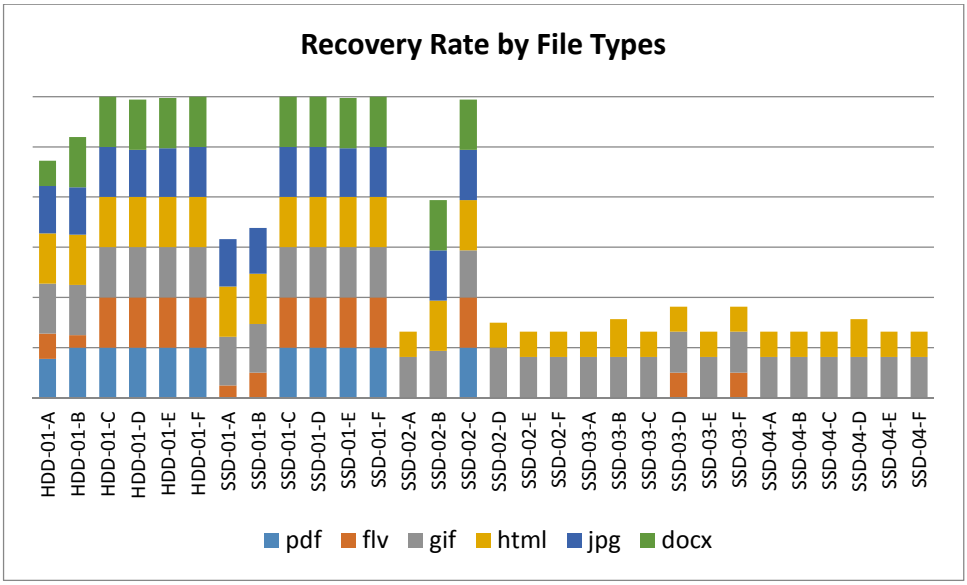

Figure 4.7: Rate of recovery for the file types displayed in stacked column chart

In total, six file types consists of 87 files were deleted for the recovery test. Table 4.26 shows the average size and average recovery rate for the deleted files for each file type. The average recovery rate was calculated by dividing the number of items recovered by the total number for each file type.

The highest recovery was seen in gif file types, which also had the smallest average file size. The second highest was recorded with html file, but the average size was the third smallest. The second smallest average size was jpg files, which had the third rate of recovery.

Table 4.26: Average size and recovery rate for deleted test files

File Types	Average of Size (Byte)	Average of Recovery Rate (%)
docx	141,748	38%
flv	4,679,074	38%
gif	1,138	90%
htm	43,315	75%
jpg	4,476	45%
pdf	182,268	36%

Further analysis was conducted and presented in Figure 4.8. Scatter plot chart was used to plot each of deleted file for size and recovery rate across the test. Here, the average recovery rate was calculated by dividing the number of recovered test cases by the total number of test cases for each file. As indicated in Figure 4.8, file size less than 1KB (1024byte) in red line demonstrate 100% recovery rate. It is therefore a smaller file size that has better chance of a successful rate of recovery, especially if the size is below 1KB threshold, the recovery was 100%.

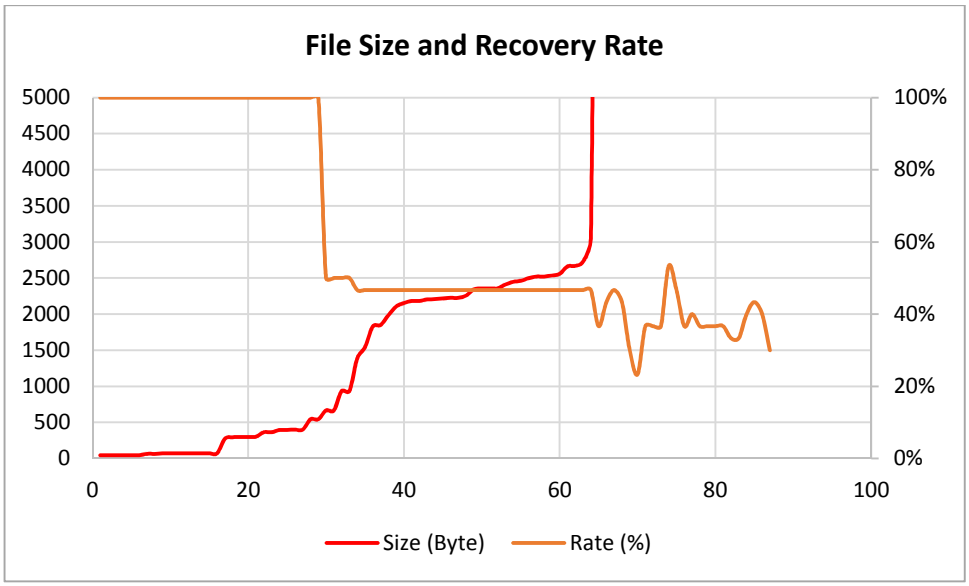

Figure 4.8: Correlation between the deleted file size (byte) and recovery rate (%)

Normalised acquisition times were plotted in Figure 4.9 to examine the overall distributions. The chart clearly segmented the results in four groups of fast and high (FH) recovery, slow but high (SH) recovery, fast but low (FL) recovery, and slow and low (SL) recovery. In general forensic perspectives, acquisition time is less of a concern as long as forensically acceptable images are obtained and the success rate of recovery is higher for the better. Therefore group FH and SH are considered desirable in generic digital forensic terms.

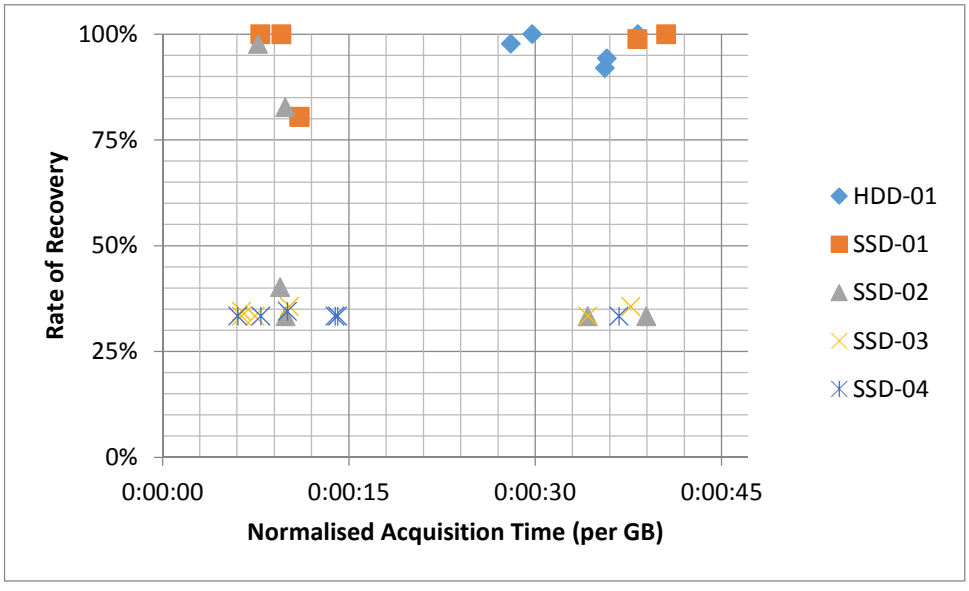

Figure 4.9: Normalised acquisition time displayed in scatter plots

HDD-01 can be found in SH group, it takes time to acquire but the recovery rate is generally high. SSD-01 provided high average recovery rate in both FH and SH groups. SSD-02 showed mixed results, which mean the success of the recovery

rate, were dependant to acquisition method. SSD-03 and 04 demonstrated consistently low recovery rates, only certain files were recovered regardless of methods and it appears acquisition time has no significant influence on these devices on the rate of recovery.

4.5 CONCLUSION

Chapter 4 has reported the variations to the methodology, laboratory testing, data collection, and analysis on the empirical results. The results from the six test cases of the four test drives were reported which included the discussion of some of the issues encountered. Visual representations of the analysis were also added to demonstrate the results graphically.

A pilot run of the proposed test was valuable, and enabled the laboratory testing to further consider details for refinement. These refinements provided the consistency, accuracy, and efficiency to the testing.

In summary, the overall results were better than anticipated based on the review of the literature and similar studies. The least recovery rate was 33% and the rate never fell below that figure. The analysis revealed that files below 1KB of size had a remarkable chance of being recovered and loaded. This speculation explained why certain files were recovered consistently. The time factor produced an inconsistent variety of results, potentially demonstrating acquisition time had less ascendancy over the chance of SSD data recovery.

Chapter 5 will discuss the results reported in Chapter 4 in relation to the research questions defined in Chapter 3.

Chapter 5

Discussion

5.0 INTRODUCTION

The literature review in Chapter 2 identified that current forensic guidelines are failing to recognise the uniqueness of SSD architecture and revolving issues experts have raised for a number of years. The critical issues SSD technology has for digital forensic investigations are technical, the definition of forensically acceptable preservation of evidence and the success rate of data recovery. These are new risks promoted by the new technology. Very little was found in the similar studies on the question of issues involved with SSD forensics and best practice guideline capability. This study set out with the aim of assessing the effectiveness of adapting the existing guidelines to deal with data stored on NAND flash media.

In this chapter, implications of these findings are discussed. The research question and sub-questions from Chapter 3 will be discussed in Section 5.1. The hypotheses are tested qualitatively in Section 5.2. The results are viewed from various angles and interpreted in Section 5.3, and potential topics for future study are introduced in Section 5.4.

5.1 THE RESEARCH QUESTION

The research question was derived in Section 3.2.2, as part of the research design. The research question is: Is it feasible to adapt existing guidelines for SSD forensics? The purpose of the research question was to verify existing forensic issues oriented with SSD, and challenge if existing guidelines are capable of handling SSD and satisfying their purpose.

The literature review in Chapter 2 provided foundation knowledge of HDD and SSD architectures and distinguished the complete differences between the two. Dissection of well-established forensic guidelines was critically examined so the unique characteristics of NAND flash technology were identified. None the guidelines mentioned the importance of risks involved with SSD. The

risks identified were inability to verify forensic copies and self-destructive behaviour permanently purging cells where files were deleted.

Section 4.4 presented the results with graphical representation and empirical analyses. The result demonstrated one out of four tested SSD recovered deleted files exceedingly well, while other three struggled and only files with size less than 1KB were consistently recoverable. The review of similar studies in Section 3.1 indicated that some files can be recovered but not loadable, therefore the result obtained were better than anticipated.

The interpretation of the results was that the chance of being able to recover deleted files successfully from SSD is below 25% depending on the manufacturer and model. The length of acquisition time seemed to have less influences on the rate of recovery but method of acquisition does. Certain combinations of SSD model and method seem to improve the rate of recovery, but the effect is insignificant.

As the literature review discussed in Section 2.2, entry level to join the SSD manufacturing business is considered low compared to HDD manufacturing. The most profitable component is the firmware software therefore attracting vast number of competitors across the globe. The situation behind the scene is chaotic, where there is no set standard with the firmware design and it is tightly protected as proprietary core component. Digital forensics has no way of reverse engineering the secrets and even if chips are extracted, reconstruction of data is near impossible.

The current guidelines are developed to effectively secure magnetic storage devices in a forensically acceptable manner, but not NAND flash devices. All the techniques involved, such as verification of software, the data recovery method, and forensic image tools, these were all developed for HDD and SSD was not covered. Although the test result was better than expected, the recovery rate was still poor in comparison with HDD. Hence, the verification issue is still not resolved which forces current SSD forensic images to be placed in a grey area when the integrity of the forensic image as evidence is questioned in a court of law. These various uncertainties between digital forensics and SSD require further clarification. For the above reasons the answer to the research question is that the existing forensic tools and techniques were not designed and capable of handling SSD, therefore it is illogical and not feasible to accept at this stage.

5.1.1 Sub-Questions

The research question was divided into four associating sub-questions in Section 3.2.2, to aid answering the research question by focusing logical reasoning onto specific issues.

The sub-question 1 was: what is a forensically acceptable SSD acquisition method? The literature review in Section 2.3.3 defined that due to the nature of SSD architecture, inability to retain original integrity is a threat to forensic investigation. Traditionally evidence integrity was verified with MD5 or SHA1 hash algorithms, and reproducibility of forensic investigation takes critical role when analysis of the evidence is challenged by any opponent. Ultimately evidence admissibility is in doubt if evidence integrity cannot be verified and leaves a potential risk of artificial tampering. Absence of a reliable data recovery method is of less concern in terms of court admissibility, but unable to prevent self-destructive behaviour while investigation is carried out on SSD is a great concern. The answer to the sub-question 1 is that SSD acquisition methods which prevent self-destructive behaviour and retain data integrity throughout the digital forensic lifecycle. Such a method is yet to be developed and hence forensic investigations are exposed to a significant threat.

The sub-question 2 was: what changes can be observed if TRIM is disabled prior to live acquisition? Details of the technique were discussed in Chapter 2; the live acquisition is rapidly becoming common approach due to its ability to capture certain information while the power is on. Such an approach requires greater knowledge and experience to determine the state of conditions surrounding the evidence and suitable decisions have to be made depending on situations. Therefore it is not suitable for every first responder. However the effect of TRIM is one of few function investigators can eliminate and must be tested. Theoretically disabling TRIM prior to acquisition should improve the chance of recovering deleted data. The result demonstrated disabling TRIM made a slight improvement in data recovery rates for most instances. The SSD-02 recovery rate doubled when TRIM was disabled. In comparison, the overall recovery rate from the live acquisition fell below the average rate from the dead acquisitions. The answer to the sub-question 2 is that when the TRIM command is disabled, the chance of data recovery dramatically improved for certain SSD models but otherwise remained unchanged. Where recovery of data is concerned,

investigations should consider the use of dead acquisition as a higher possibility of recovering data. But as always, some trade-offs are involved when critical decisions have to be made.

The sub-question 3 was: What is the most effective SSD friendly imaging method? The term effective suggests as complete as possible so the data recovery is possible from the acquired image. The purpose of this question was to establish if a certain combination of existing acquisition methods can improve the rate of data recovery on a SSD. As discussed in Section 4.4.2, statistically the method used in test case C had the highest rate of recovery amongst others. The method used Tableau TD2 hardware duplicator for the acquisition and EnCase was used for the deleted file recovery. However the outcome was not consistent across the test SSDs. The method was demonstrated only to be applicable to certain drives, but not all of them. The determination of which drive model certain acquisition methods can be applied is unknown. Even when such a determination technique is developed, selecting correct device and methods out from a considerable number of combinations is impractical. Therefore the answer to the sub-question 3 is that the effect method of acquisition varies depending on SSD manufacture and model, and some models have harsher built-in sanitising processes which permanently destroy the residual data.

The sub-question 4 was: how much more data can be recovered from SSD with the compound forensic collection procedure? The review of similar studies in Section 3.2.1 indicated that recovery of loadable files were extremely difficult without directly reading off extracted NAND chips. Often even files were recovered blank, actual content of the file remained unrecovered. The results in Section 4.4 demonstrated minimum of 33% of delete files, or files under 1KB have significantly increased chance of being recovered and remain loadable. Therefore the answer to the sub-question 4 is that when the compound forensic collection procedures are followed as a guideline, smaller files have higher chance of recovery and larger files. All be recovered up to 100% if certain conditions are satisfied. But a higher rate of recovery remains reliant on the unpredictable automated SSD firmware manoeuvre which widely varies depending on manufacturer and device models.

5.2 HYPOTHESES TESTING

Four hypotheses developed in Section 3.2.3 will be examined and validated against the test case results. Table 5.1 – 5.4 display the hypotheses associated with arguments and state conclusions drawn respectively.

5.2.1 H1: Current forensic guidelines are incapable of handling SSD

The results support the assumption H1 made is correct, despite two test cases demonstrated 100% data recovery. Acquisition and recovery methods used in the test cases were not developed for SSD and various uncertainties especially around the firmware prohibiting existing guidelines were unpredictable. Therefore the adaptation of the guidelines for SSD forensics is unreliable and incapable.

Table 5.1: H1 testing

For	Against
Acquisition and recovery technique are developed for magnetic storage, not NAND flash storage.	Two test cases proven that certain combinations of acquisition method allows up to 100% data recovery.
Integrity of acquired image can't be verified which makes the evidence vulnerable to tampering.	It is not only SSD that integrity of original data cannot be verified.
Recovery rate of deleted data is considerably reduced compared to HDD.	
Architectural diversities are not distinguished and incorporated into the existing guideline.	
Conclusion: Accept	

5.2.2 H2: Data recovered from SSD is not loadable

The review of similar studies revealed data recovery with SSD is troublesome and often files are recovered in a non-loadable state.

Table 5.2: H2 testing

For	Against
File size above 1KB struggled to be recovered especially in loadable condition.	Various file types and sizes were successfully recovered in loadable condition.
	Rate of recovery was measured when recovered files were confirmed loadable, and all the test cases presented 33% minimum rate of recovery.
Conclusion: Reject	

The test cases simulated a situation where files are deleted from a computer then investigation attempted to recover them. The result showed various responses depending on combinations of SSD model and method. However files were recovered in loadable conditions especially if the size was below 1KB, therefore H2 is falsified.

5.2.3 H3: Faster imaging method provides better rate of data recovery

Theoretical analysis indicated that because certain built-in processes are automated there is no way of disabling them during the acquisition, faster acquisition speed should minimise the chance of self-destruction process. However the results demonstrated the longest acquisition could still achieve near complete data recovery. The results are demonstrating mixed responses primarily based on certain combinations and regardless of acquisition time. Nevertheless, it is clear that the presence of the self-destructive processes cannot be lightly ignored therefore the test for H3 is inconclusive.

Table 5.3: H3 testing

For	Against
Five test cases had fast and high recovery rates while five test cases had slow and low recovery rates.	Three test cases had slow but high recovery rates, and nine test cases had fast but low recovery rates.
Theoretically longer acquisition time allows more opportunity for the firmware to sterilise the deleted storage space.	The longest acquisition time recorded in the experiment had 99% recovery rate.
Conclusion: Inconclusive	

5.2.4 H4: Certain combination provides better rate of data recovery

The key to SSD performance is solely governed by its firmware, and the commercial competition is widely welcoming any new comers willing to participate in the new storage technology race. This situation has the market flooded with various firmware distributions; it is common to find the same manufacturer adapting multiple firmwares for every model they release without consistency or any standard. Hence the test results demonstrating the same method produces variety of results depending on SSD models. In one case, it is verified that a certain combination allowed the investigation to recover 80% of the data, while another method only had a 33% recovery rate. Therefore the results clearly supports that assumption made in H4 is true.

Table 5.4: H4 testing

For	Against
The rate of recovery varied depending combinations of SSD and acquisition method.	Not all the test cases were repeatedly tested multiple times to verify the results.
One test case repeatedly demonstrated exceedingly high recovery rates for only one specific method.	
Conclusion: Accept	

5.3 DISCUSSION OF FINDINGS

The result of this study indicated that deleted data can be recovered in a loadable state for up to 100% from SSD, but the possibility is hard to predict. The only logical pattern detected from the results was that files under 1KB size had a significantly high chance of successful recovery, while larger files were statistically less successful. The recovery rates were inconsistent depending on the combination of SSD model and acquisition method. Disabling the TRIM command demonstrated a better rate of recovery in one test case, but unable to detect its effect from other cases. There are large amounts of uncertainty associated with the application of current forensic practice to data stored on NAND flash media, and the results of this experiment found that the guidelines are not suited to effectively handle SSD media.

This study supports Bell and Boddington's (2010) statement which described the application of existing evidence collection processes for SSD as imprudent and potentially reckless, assumptions about the behaviour of storage media is no longer valid, and such practice frustrates post recovery forensic analysis.

The findings also collaborate with Bednar and Katos's (2011) research, which stated that the information and advice contained in ACPO guideline cannot be used for handling of SSD, especially the use of live acquisition techniques was not recommended due to the presence of a garbage collection process which purges the residual data during the acquisition. Their statement is in alignment with the findings from this study, live acquisition with or without TRIM produced a less successful rate of recovery. Bednar and Katos's findings must be interpreted with additional cautions because live acquisition has another main purpose for

usage and critical decisions have to be made that can give situations where recoverable data must be sacrificed. That is regardless of if SSD or HDD was the source drive.

The results marginally contradict King and Vidas's (2011) conclusion. In their study no data was recoverable when TRIM was enabled, and a high recovery rate was obtained with TRIM disabled. This experiment identified only one test case where the effect was disabling TRIM lead the rate to improve. In addition, the Intel drive had the lowest recovery rate in their research while one of Intel drive (SSD-01) demonstrated unusually high rate of recovery across all the test cases. However both results show that the statistics appear to be manufacture dependant, which directly supports the assumption made in H4 of a certain combination can provide an improved recovery rate. The firmware process is suspected as a cause of this unpredictable phenomenon, but the experiment did not provide any evidence.

5.3.1 Evaluation of Testing Methodology

The literature review in Chapter 2 identified some of the technical limitations in relation to SSD forensics. The issues consist of three interlinking components of unavailability of SSD firmware algorithms, premature development of SSD forensic techniques and tools and poor distribution of risks involved with SSD.

The review of similar studies in Section 3.2.1 examined the specific tools and techniques for SSD showing that it may take time to develop. There is no method available to create a completely duplicate image of SSD in forensically acceptable manner. However none of the trusted guidelines include risks involved with SSD is an serious issue especially for first responders dealing with evidence at the scene. While tools and techniques are under urgent research and development, the guidelines must at least educate the risk widely raising awareness and concerns amongst everyone involved in the investigation. At the same time, the best possible method had to be found and to identify what existing practice or combinations of tools and techniques are still effective for SSD.

In Section 2.4.3, six commonly referenced best practices and guidelines were critically examined and contrasted to identify gaps in the forensic processes, especially emphasising the preservation of evidence and data recovery analysis. The gaps identified in the similar studies and guidelines were combined together

and formed the compound guideline aiming to maximise the chance of data recovery when a SSD was involved in an investigation.

A test case experiment was designed and processes of interests were tested for effectiveness. The effectiveness was measured empirically by comparing a number of test files recovered in loadable condition following the six test cases. The result from HDD was used as value of control, provided benchmark recovery rate for each test case. The term recovered was defined as a file restored in a loadable condition. The results exceeded expectation and files under 1KB were consistently recovered across the test cases. Especially the recovery rates from SSD-01 were unexpected.

5.3.2 High Recovery Rate

As Section 4.4.1 presented the results grouped by test drives, the rates SSD-01 scored were almost equal to the benchmark figure. 80% recovery rates where live acquisition methods, near 100% average recovery rates were scored with remaining dead acquisition tests. These were overly beyond acceptable figures distorting the overall outcome and require additional logical reasoning for explanation. Firstly human error was suspected but analysis on test logs shows all procedures were followed correctly. If human error was involved, a similar error is expected from other results as well. The only logical explanation this research could determine was the device itself. It is uncertain that if the firmware or the NAND flash was the root cause, but the experiment was specifically designed and controlled only the SSDs are the sole variable and therefore a variation in the result can only occur when the test drive responded differently. However the experiment was not designed to identify what causes the different results, it was designed to identify the acquisition method which maximised the chance of data recovery.

5.3.3 Deletion to Acquisition Interval

The timing of the acquisition maybe responsible for the higher than expected recovery rates. All the test cases were given five minutes interval between the time of deletion and acquisition. In theory, the longer the interval promotes less chance of recovery because of the TRIM and garbage collection effects purging residual data from the deleted space. Real investigation rarely has opportunity to acquire storage device within five minutes after deletion, but at the same time it is

hard to estimate the average time between deletion and acquisition. In real life, further complications can be expected, such as use of secure delete tools, encryption, or even steganography. However the primal focus of this research is to empirically examine current forensic technique to seek to maximise chance of data recovery. It is expected the rate will diminish as the interval is stretched longer, but it would be interesting to see if the threshold file size for higher rate of recovery will also diminish.

5.3.4 Acquisition Speed

There is no method of interrupting SSD firmware processes unless NAND flash chips are extracted from the PCB. This allows the firmware to be active during the course of acquisition hence the original data structures become dynamic and disallow investigation to retain a hash verified forensic image. The longer acquisition takes the firmware could initiate destruction processes as a result, less recoverable data will be available for the investigation. For this reason, H3 assumed that a longer acquisition time on SSD will diminish the success rate of recovery. The result did not support the hypothesis. The longest acquisition time recorded was a test case SSD-01-F and the recovery rate was 100%. As discussed previously in this section, the results obtained from SSD-01 are peculiar, further analysis without SSD-01's result may assist with logical reasoning. However the results remained segmented and no sign of declined or diminishing rate of recovery was discovered. As presented in Section 4.4.3, the recovery rate and acquisition time had four segmented relationships. When the result of SSD-01 was excluded, SH (slow but high recovery) group was eliminated but the others remained in the same three segmentations. The results demonstrated contradicting outcomes between the critical literature reviewed and the field testing and remain inconclusive.

5.3.5 Sample Size and Selection

The sample size in this research was discussed as a limitation in Chapter 3. As the research's empirical analysis progressed, selection processes and the available number of sample drives became a minor concern. Especially when the result demonstrated that the recovery rates were dependant on a combination of certain models and methods, and additional factors of SSD became important for further specific analysis.

The number of bits stored in each cell is define as SLC (single), MLC (multi), or TLC (triple) and increase in bit per cell means higher cost performance. The experiment used one TLC drive (SSD-02), and remaining was all MLC. SLC was not included. Storage capacity, read speed, and DRAM cache size were not the same and its influence is uncertain but cannot be underestimated. However these factors did not deteriorate the quality of this research and beyond the scope set out in the research design.

5.4 SUGGESTIONS FOR FUTURE WORK

This research identified that existing guidelines are ignoring the fact SSD is developed based on a completely different scientific foundation. The experiment also confirmed that available acquisition methods developed for HDD is not valid on SSD, various uncertainties surrounding SSD technology became a threat to forensic investigation. The following sections will discuss the potential field of future research. Section 5.4.1 will discuss that the ultimate forensic solution requires extensive research and development. Section 5.4.2 will discuss the options without new invention. Section 5.4.3 will discuss the ideal solution which solves the issues.

5.4.1 SSD Root Toolkit

It is apparent further work is desperately required to aid this crisis and develop a new guideline which first responders and investigation can comfortably reference to and rely on. Ultimately development of SSD firmware root access toolkit will solve many issues. Manufacturers are reluctant to disclose specifications hidden in the firmware, but if custom firmware can be developed and overwritten using firmware update protocols, then full control of SSD operations is surrendered to investigators.

5.4.2 Improved Admissibility

Until such new methods become available, it would be useful if the method of maximising admissibility on unverified forensic images can be studied. Investigations require knowledge on how unverified forensic images can maximise the quality of admissibility in a court of law. Such knowledge is also transferrable to other fields of acquisition such as cloud forensics and volatile data analysis. It is ironic that NAND flash has conquered the volatility issue once, but

consumer's eagerness for better performance leave SSD's residual data volatile again.

5.4.3 Implementing Read-only Mode

Alternatively, it would be interesting to study the manufacturer vision with SSD's long term development. It is matter of time that general consumers become aware of recovery issues involved with SSD and if one manufacturer could implement a dip switch or a jumper to enable read-only mode to preserve the state, most the issues discussed in this research will evaporate.

5.5 CONCLUSION

This chapter explored the results reported in Chapter 4, then evaluated the implications of the findings. The research question and sub-questions were answered in terms of evidence provided in Chapter 4. There is no single solution to data recovery for SSD.

Qualitative analysis was used for the hypotheses and the conclusions were stated. The findings were referenced to the literature review in Chapter 2 and the research design in Chapter 3 for benchmarking and evaluation, and then the discussions progressed for logical reasoning and logical explanations. Overall the findings were supportive and in alignment with the similar studies reviewed in Chapter 3.

Three topics were provided as suggestions for the future work. Each topic has its own value but also comes with different challenges. In the next chapter, conclusions including reflection on this research, significance of the findings, and recommendations for future studies will be presented.

Chapter 6

Conclusion

6.0 INTRODUCTION

This research has evaluated existing computer forensic guidelines for the suitability of handling SSD. The research findings supported the hypothesis and none of the existing guidelines are fully capable for handling SSD at the present time. In Chapter 6, a belief based on the logical reasoning and evidence the current study has accumulated will be presented.

A brief summary of this research is presented in Section 6.1, followed by restatement of objectives in Section 6.2. Section 6.3 to 6.5 covers concise statements of the research findings, suggesting implications, and significance of the findings. The limitations are revised in Section 6.6 and associated recommendations for future research are presented in Section 6.7.

6.1 SUMMARISING THE CONTENT

This research has examined the advent of SSD technology and issues brought to computer forensic investigations. The literature review explored the evolution of the main data storage used in ordinary computers. HDD is the most common computer storage device for decades, now SSD is slowly taking over its position and manufacturers are developing and marketing their new models at a high rate. The problems arising from SSD for digital forensics are not widely discussed and even forensic guidelines are not making these issues obvious to alert investigators and others to the new challenges forensic investigations are facing at the moment.

The experiment was designed to measure the guideline's capability of handling SSD. Instead of fitness testing the six guidelines, each guideline was subdivided into processes and qualitatively assessed for SSD forensic readiness. The further review on the similar studies was conducted and together with the relevant processes, the compound guideline to maximise chance of data recovery from SSD was developed. The research sub-questions were derived from the guideline processes to structure the experiment, and then the data map was drawn to visualise the flow of information.

The results were unexpected but satisfactory and generally supported the hypotheses derived from the literature review. The current acquisition methods are less effective with SSD but certain combinations allow a higher rate of recovery. However the logic behind the successful combination is full of uncertainties and results are unpredictable. It appears the root cause reside within the SSD firmware, the manufacturer trade secret black-box. The firmware holds the key algorithm causing evidence destruction behaviour and because there are numerous firmware variations flooding the SSD market without any form of standards, each SSD behaves different to another.

The discussion proposed options based on the research findings and all the proposals have one thing in common. The SSD firmware must be suppressed for successful preservation of evidence and recovery of data.

6.2 SUMMARISING THE FINDINGS

The most obvious finding to emerge from this study is that complete recovery of deleted files on a SSD is not impossible, but the possibility varies depending on the acquisition timing and the certain combination of SSD models and acquisition methods. Pre-determination of the best combination is difficult, and requires better understanding of the black-boxed SSD firmware that manufacturers will not easily disclose. Extensive data recovery tests on every SSD model and acquisition methods may provide a growing list of statistical probability, but such an approach is resource intensive and will never solve the uncertainties.

The second major finding was that smaller files can be recovered better. The result demonstrated that especially files with less than 1KB had significantly increased chance of recovery. It is uncertain why the threshold is 1KB, but popular file types such as .pdf, .docx, .xlsx, and .jpg are generally over 1KB on modern computers. Multimedia files and mailbox files are even bigger and finding of higher recovery rate with smaller files does not help the situation at all.

It was also shown that disabling TRIM process is effective to a limited extent. Disabling TRIM does not stop the firmware's destructive behaviour. But if live acquisition for volatile data is in high priority, then it is recommended to disable the TRIM to eliminate potential. The results indicated that one out of four drives showed better recovery rate when TRIM was disabled.

6.3 SUGGESTING IMPLICATIONS

The findings from this study suggest that the growing number of SSD users have become a new threat to the computer forensics investigations. It is unknown why SSD was treated like HDD to start with. It is obvious SSD is designed to act like HDD, only to perform better than HDD. But its architecture, components, and method of retaining data are all dissimilar only the physical size and connector interface follows the same standard. This study revealed the current guidelines, acquisition methods, data recovery tools and techniques are all developed based on magnetic storage devices. The reason is because there is nothing else to recover the data from for computer related evidences. Historically, till recently users had no technology to retain data stored on RAM without a continuous supply of power. As soon as power is off, data on the RAM vanished hence the reason it was called volatile memory. The new era has come and technology allows electrons to be trapped without a power supply. But tools and techniques developed for magnetic storage are not the root cause of the issue.

The issue is not where the data is stored at but how the data is stored and administrated. It is the manufacturer's solution to the consumer demand for the better performance at an affordable price. Data stored on a SSD are compressed, encrypted, and residual information is periodically purged without interference from users. Effectively optimised commercial solutions embedded into SSD firmware is the root cause of the unresolved situation. And computer forensics isn't the only group affected.

6.4 SIGNIFICANCE OF THE FINDINGS

The findings from this study make several contributions to the current literature. Firstly the present study provides additional evidence with respect to SSD forensics and state of the existing computer forensics guidelines. As discussed in Chapter 2, all six guidelines are resourceful and well reputed. Each of them is composed diversely for a different target audience and their own purposes as well. However they all significantly lack awareness of SSD technology complexities.

The methods used to determine SSD forensic readiness for the guidelines in Section 2.4.3 will serve as a base for future studies and understanding of the slightly different role of each guideline. Some are at higher level of scope and less focused on technical aspects while some are highly technical and more suitable

for skilled experts. The workflow charts are also helpful to visually understand what processes are covered in the guideline. The charts were drawn for the same generic purpose therefore it is easier for qualitative analysis as well.

As mentioned in Section 6.3, data recovery from SSD is possible but number of uncertain variables, especially the black-boxed SSD firmware holds the key to progress further. The computer forensics society isn't the only group interested in advanced data recovery from SSD. It is matter of time for the general consumers to widely realise the nature of inability to recover data from SSD once deleted, and once they do the manufacturers will be lobbied to arrange some sort alternatives to rectify the issue. The easiest is to provide a switch (preferably hardware) to enable read-only mode on SSD. The switch will allow users to select preservation of the state of SSD, including residual data, and enables HDD data recovery tools and techniques to be effective once again.

6.5 LIMITATIONS OF THE RESEARCH

A number of important limitations to this research need to be considered. First, the sample size was relatively small. The literature review revealed that there are numerous manufacturers competing in SSD market, the lower entry requirement and high profitability attracted various new vendors to join the competition. Only four SSDs were tested in this research and generalisation of the result may require more variety of SSD devices.

The study did not run the same test more than once. There were two occasions the recovery rates were unexpectedly high and re-tested to verify the accuracy. The results were verified correct. Due to the small sample size used in the research, multiple test runs might have produced additional supportive statistics.

Thirdly, consistency of tested SSD specifications is another limitation to be considered. The experiment focused on the effectiveness of acquisition methods and the effectiveness was measured by the rate of successful data recovery. The test SSD capacity size ranged from 120GB to 256GB. The variations in storage capacity impacted the duration time for acquisition and normalisation was required when acquisition time was compared. In theory, a larger volume requires a longer time to acquire and therefore a higher risk of exposure to the firmware destruction processes. Correlation between acquisition

time and rate of recovery was part of hypotheses tested in Section 5.2.3. However the result didn't show a negative correlated relationship between the acquisition time and rate of recovery. The quantity of bits per cell was also outside the research scope, MLC and TLC were used in the testing and SLC was not involved. The only requirement was availability of TRIM command support which all four test SSDs supported.

In relation to time factor, time interval was set exactly five minutes between the time of deletion and beginning of the test acquisition processes. In the similar study, Bell and Boddington (2011) used 15-minute interval which resulted with substantial garbage collection erasure. The variable in the time interval was outside the scope of this research.

6.6 RECOMMENDATIONS FOR FURTHER RESEARCH

This research has thrown up many questions in need of further investigation. Manufacturers are keen to develop more cost effective SSD by data compression and higher density bits per cell architecture. Further study in data recovery and higher density bits per cell correlation will be beneficial.

Research in direct NAND chip access without physical extraction will provide alternative approach to bypass the firmware, but at the same time reverse engineering built-in encryption and compression algorithm is essential.

Another interesting field of research is to utilise SSD firmware update process to hijack the SSD controller by uploading forensic custom firmware. Since manufacturers are reluctant to reveal their firmware algorithm, a substitute is required to at least preserve the evidence and allow read-only access to SSD storage area.

Alternatively, research in the importance of having read-only switch on storage device may assist manufacturers to consider implementing such option, either hardware dip switch or software, allowing users to manually write-block the storage device to allow maximise chance of data recovery.

Overall, it is obvious that further collaborative research work with SSD manufacturers are the most recommended and proactive approach to mediate current disastrous situation. Ultimately, manufacturers must develop an industry standard forcing all storage devices to be equipped with options to enhance the

chance of data recovery if required. In the meantime, the forensic guidelines must incorporate the issues and options revolving around SSD storage device.

References

ACPO (2010, March 26). *"Our Structure"* Retrieved August 18, 2010, from www.acpo.police.uk/about_pages/structure.html

Aerocool (2012, May 5). Superiority of FDB [Diagram]. Retrieved from http://www.aerocool.us/accessory/shark_b12.html

Ashcroft, J., Daniels, D., & Hart, S. (2004). *Forensic Examination of Digital Evidence: A Guide for Law Enforcement, April 2004* (NCJ 199408). Retrieved from Office of Justice Programs National Institute of Justice website: http://www.ojp.usdoj.gov/nij

Ballou, S., & Gilliland, R. G. (2011). Emerging paper standards in computer forensics. *Digital Investigation*, *8*(2), (pp. 96-97). doi:10.1016/j.diin.2011.05.017

Barker, E., Barker, W., Burr, W., Polk, W., & Smid, M. (2012). *Recommendation for key management: Part 1* (NIST 800-57). Gaithersburg, MD: U.S. Dept. of Commerce, National Institute of Standards and Technology.

Bednar, P. M., & Katos, V. (2011). SSD: New Challenges for Digital Forensics. *VIII Conference of the Italian Chapter of AIS Information Systems: a crossroads for Organization, Management, Accounting and Engineering Proceedings ISBN: 978-88-6105-063-1*.

Bell, G., & 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*, *5*(3), (pp. 1-20). Retrieved from http://www.jdfsl.org/subscriptions/abstracts/abstract-v5n3-bell.htm

Bestofmedia (2011, August 30). *Read/Write Head Designs: Ferrite, Metal-In-Gap, And Thin-Film - Hard Drives 101: Magnetic Storage*. Retrieved April 30,

2013, from http://www.tomshardware.com/reviews/hard-drive-magnetic-storage-hdd,3005-2.html

Blattau, N., & Hillman, C. (2004). Failure Mechanisms in Electronic Products at High Altitudes. *CALCE Electronic Products and Systems Center*. Retrieved from http://www.dfrsolutions.com/pdfs/2004_HighAltitude_Hillman-Blattau.pdf

Blount, W. C. (2007). *Why 7200 RPM Mobile Hard Disk Drives?* Retrieved from Hitachi Global Storage Technologies website: http://www.hgst.com/

Bogan, C. E., & English, M. J. (1994). *Benchmarking for best practices: Winning through innovative adaptation*. New York: McGraw-Hill.

Britz, M. (2009). *Computer forensics and cybercrime: An introduction*. Upper Saddle River, N.J: Pearson Prentice Hall.

Brooker, W. (2005, March 10). *Toshiba MK1031GAS 100GB 2.5" Hard Disk*. Retrieved August 10, 2013, from http://www.3dvelocity.com/reviews/toshiba/mk1031gas_4.htm

Bruker (2011, June 1). *Magnetic Force Microscopy – MFM » Bruker Blog*. Retrieved July 1, 2013, from http://blog.brukerafmprobes.com/2011/06/magnetic-force-microscopy-mfm/

Carrier, B. (2011). *File system forensic analysis*. Upper Saddle River, NJ [u.a.: Addison-Wesley.

Chen, C. J. (1993). *Introduction to scanning tunneling microscopy*. New York: Oxford University Press.

Christensson, P. (2005). Cluster Definition. In *The Tech Terms Computer Dictionary*. Retrieved March 3, 2013, from http://www.techterms.com/definition/cluster

Clausing, J. (2009). Fight crime Unravel incidents one byte at a time. *SANS Computer Forensics and e-Discovery*. Retrieved from http://memory8.com/fight-crime-unravel-incidents-one-byte-at-a-time-w1446/

Cooke, J. (2011). *NAND 201: An Update on the Continued Evolution of NAND Flash* . Retrieved from http://www.micron.com/~/media/Documents/Products/White Paper/nand_201.pdf

Coughlin, T. M. (2008). *Digital storage in consumer electronics: The essential guide*. Amsterdam: Elsevier/Newnes.

Dpi, A. (2008, September 25). *Hard Disk Drives*. Retrieved March 9, 2013, from https://www.itschool.gov.in/PDF/SITC%20hardware%20training/Hard%20disk.pdf

Duffy , D. (2004, May 1). *Computer Forensics Investigations: Body of Evidence - CSO Online - Security and Risk*. Retrieved May 3, 2013, from http://www.csoonline.com/article/219226/computer-forensics-investigations-body-of-evidence?page=1

Fellows, R. (2007, September 7). *Evaluating SAS Technology for Business Requirements | Serial Storage Wire*. Retrieved May 4, 2013, from http://www.scsita.org/serial-storage-wire/2007/09/evaluating-sas-technology-for-business-requirements.html

Freeman, M., & Woodward, A. (2009). Secure State Deletion: Testing the efficacy and integrity of secure deletion tools on Solid State Drives. *Proceedings of the 7th Australian Digital Forensics Conference*.

Garfinkel, S., Malan, D., Dubec, K., Stevens, C., & Pham, C. (2006). *Advances in digital forensics II: IFIP International Conference on Digital Forensics, National Center for Forensic Science, Orlando, Florida, January 29-February 1, 2006*. New York: Springer.

Garfinkel, S. L. (2010). *Building Realistic Forensic Corpora to Enable Undergraduate Education and Research* (2010-07-27 ECC 1). Retrieved from NPS website: http://simson.net/ref/2010/2010-07-27%20ECC%201%20-%20Keynote.pdf

Goh, T. B., Li, Z., Chen, B. M., Lee, T. H., & Huang, T. (2001). Design and Implementation of a Hard Disk Drive Servo System Using Robust and Perfect Tracking Approach. *IEEE Transactions on Control Systems Technology*, *9*(2), (pp. 221-233). doi:10.1.1.139.5512

Guidance (2008, August 14). *User's Guide - Tableau*. Retrieved May 2, 2013, from http://Tableau_TD2_Users_Guide.pdf

Gutienez, C. M., & Jeffrey, W. (2006). *Guidelines for media sanitization: Recommendations of the National Institute of Standards and Technology* (SP 800-88). Retrieved from U.S. Dept. of Commerce, National Institute of Standards and Technology website: http://csrc.nist.gov/publications/nistpubs/800-88/NISTSP800-88_with-errata.pdf

Gutmann, P. (1996). Secure Deletion of Data from Magnetic and Solid-State Memory. *SSYM'96 Proceedings of the 6th conference on USENIX Security Symposium, Focusing on Applications of Cryptography*, *6*, 8.

Harris, D. (2007). *NAND Flash Memory | Digital ICs content from Electronic Design*. Retrieved from Penton website: http://electronicdesign.com/digital-ics/nand-flash-memory

Hashimoto, H., Ochiai, M., & Sunami, Y. (2012). Robust Optimum Design of Fluid Dynamic Bearing for Hard Disk Drive Spindle Motors. *J. Tribol*, *134*(4). doi:10.1115/1.4007246

Herth, R. (2011, March). *SSD-Practical Solutions*. Retrieved April 9, 2013, from http://www.msc.de/en/news/pressroom/essays/7480-www/6318-www.html

Howe, T. (2008, May 6). *IBM 305 RAMAC- The First Computer with a Hard Disk Drive in 1956*. Retrieved April 29, 2013, from http://www.cedmagic.com/history/ibm-305-ramac.html

Hu, X. Y., & Haas, R. (2010). *The Fundamental Limit of Flash Random Write Performance: Understanding, Analysis and Performance Modelling* (RZ3771). Retrieved from IBM Research website: http://domino.watson.ibm.com/library/cyberdig.nsf/papers/50A84DF88D540735852576F5004C2558

Hu, X., Eleftheriou, E., Haas, R., Iliadis, I., & Pletka, R. (2009). Write amplification analysis in flash-based solid state drives. *IBM Zurich Research Laboratory*. doi:10.1145/1534530.1534544

Hutcheson, G. D. (2009). The Economic Implications of Moore's Law. *Into The Nano Era, Springer Series in Materials Science*, *106*, (pp. 11). doi:10.1007/978-3-540-74559-4_2

IBM (2001, September 14). *Meet Your Hard Drive*. Retrieved April 20, 2013, from http://www.research.ibm.com/research/gmr/basics.html

Jang, E., Koh, B., & Choi, Y. (2012). A study on block-based recovery of damaged digital forensic evidence image. *Multimedia Tools and Applications*, *57*(2), (pp. 407-422). doi:10.1007/s11042-011-0738-9

Jung, M., & Kandemir, M. (2013). Revisiting widely held SSD expectations and rethinking system-level implications. *SIGMETRICS Perform*, *41*(1), (pp. 203-216). Retrieved from doi:10.1145/2494232.2465548

Kawaguchi, A., Nishioka, S., & Motoda, H. (1995). A FlashMemory Based File System. *TCON'95 Proceedings of the USENIX 1995 Technical Conference Proceedings*, 13. doi:10.1.1.41.6774

Karbo, M. (2011). Chapter 43. SCSI, USB and Firewire. In *PC Architecture*. Denmark: ELI Aps.

Kent, K. A., Checalier, S., Grance, T., & Dang, H. (2006). *Guide to integrating forensic techniques into incident response: Recommendations of the National Institute of Standards and Technology* (NIST SP800-86). Retrieved from U.S. Dept. of Commerce, Technology Administration, National Institute of Standards and Technology website: csrc.nist.gov/publications/nistpubs/800-86/SP800-86.pdf

King, C., & Vidas, T. (2011). Empirical analysis of solid state disk data retention when used with contemporary operating systems. *Digital Investigation*, *8*, (pp. 111-117). doi:10.1016/j.diin.2011.05.013

Kioskea (2013, January). *Formatting - Formatting a hard drive*. Retrieved April 22, 2013, from http://en.kioskea.net/contents/626-formatting-formatting-a-hard-drive

Kissel, R., Skolochenko, M., & Li, S. (2006). *Guidelines for media sanitization: Recommendations of the National Institute of Standards and Technology*

(SP800-88). Retrieved from U.S. Dept. of Commerce, National Institute of Standards and Technology website: http://csrc.nist.gov/publications/nistpubs/800-88/NISTSP800-88_with-errata.pdf

Kitagawa, T. (2011, September 10). *DOS/V POWER REPORT - SSD*. Retrieved June 3, 2013, from http://www.dosv.jp/other/0910/02.htm

Kleinert, J., & Leitner, A. (2008, January). Friend of Flash - Flash memory and the LogFS filesystem. *Linux Magazine*, (86), (pp.40-41). Retrieved from http://www.linux-magazine.com/w3/issue/86/040-041_logfs.pdf

Kozierok, C. M. (2001, April 17). *PCGuide - Ref - Hard Disk Operational Overview*. Retrieved May 1, 2013, from http://www.pcguide.com/ref/hdd/op/over.htm

Lee, S. W., Park, S. J., Campbell, E. E., & Park, Y. W. (2011). A fast and low-power microelectromechanical system-based non-volatile memory device. *Nature Communications*, *2*(1227), 1. Retrieved from doi:10.1038/ncomms1227

Lyle, J. R. (2003). NIST CFTT: Testing Disk Imaging Tools. *International Journal of Digital Evidence*, *1*(4), 1. Retrieved from http://www.utica.edu/academic/institutes/ecii/publications/articles/A04BC142-F4C3-EB2B-462CCC0C887B3CBE.pdf

Mackenzie, D., Meyering, J., Paterson, R., Pinard, F., Berry, K., Youmans, B., & Stallman, R. (2009). *GNU Coreutils Manual* (8.22). Retrieved from Free Software Foundation, Inc website: http://www.gnu.org/software/coreutils/manual/coreutils.pdf

Marshall, A. M. (2009). *Digital Forensics: Digital Evidence in Criminal Investigations*. Chichester: John Wiley & Sons.

MediaSmarts (2013, April 15). *Deconstructing Web Pages*. Retrieved May 19, 2013, from http://mediasmarts.ca/sites/default/files/pdfs/lesson-plan/Lesson_Deconstructing_Web_Pages.pdf

Medlin, B. D., & Crazier, J. A. (2010). A Study of Hard Drive Forensics on Consumers' PCs: Data Recovery and Exploitation. *Journal of Management Policy and Practice*, *12*(1), 27. Retrieved from http://www.na-businesspress.com/JMPP/MedlinWeb.pdf

Micheloni, R., & Eshghi, K. (2013). SSD Architecture and PCI Express Interface. *Springer Series in Advanced Microelectronics*, *37*, (pp. 19-45). doi:10.1007/978-94-007-5146-0 2

Mueller, S., & Micro House (Firm) (1998). *Micro House PC hardware library*. Indianapolis, Ind: Que.

Mueller, S. (1999). *Upgrading and repairing PCs*. Indianapolis, Ind: Que.

Mueller, J. (2006, April 19). *STM References - Annotated Links for Scanning Tunneling Microscope Amateurs*. Retrieved April 20, 2013, from http://www.e-basteln.de/index_r.htm

NCJRS (1998). *Evolution and Development of Police Technology* (NCJ 173179). Retrieved from SeaSkate, Inc website: https://www.ncjrs.gov/App/publications/abstract.aspx?ID=173179

Nguyen, M. (2013). *Samsung 840 Pro Review (256GB) | Ubergizmo*. Retrieved from ubergizmo.com website: http://www.ubergizmo.com/2013/05/samsung-840-pro-review-256gb/

Nisbet, A., Lawrence, S., & Ruff, M. (2013). A Forensic Analysis and Comparison Of Solid. State Drive Data Retention with Trim Enabled File Systems. *Proceedings of the 11th Australian Digital Forensics Conference*, (pp. 103-111). Retrieved from http://ro.ecu.edu.au/adf/125/

NIST (2002). *Secure hash standard* (FIPS 180-2). Retrieved from Information Technology Laboratory, National Institute of Standards and Technology, U.S. Dept. of Commerce, Technology Administration website: http://csrc.nist.gov/publications/fips/fips180-2/fips180-2withchangenotice.pdf

NIST (2008, December 24). *NIST General Information*. Retrieved December 20, 2013, from http://www.nist.gov/public_affairs/general_information.cfm

Nolan, R., O'Sullivan, C., Branson, J., & Waits, C. (2005). *First Responders Guide to Computer Forensics* (CMU/SEI-2005-HB-001). Retrieved from CERT Training and Education website: www.cert.org/archive/pdf/FRGCF_v1.3.pdf

Nolan, Richard (2005). *First responders guide to computer forensics*. Pittsburgh, Pa: Carnegie Mellon University Software Engineering Institute.

Oura, K., & Lifshits, V. G. (2003). *Surface science: An introduction.* Berlin: Springer-Verlag.

PCstats (2006, January 11). SATA and IDE HDD connectors compared. Retrieved from http://www.pcstats.com/articleview.cfm?articleid=1778&page=2

Pearsall, J. (Ed.). (2013). Forensics. In *Oxford Dictionaries Online*. Retrieved March 1, 2013, from

http://www.oxforddictionaries.com/definition/english/forensic?q=forensics#forensic__62

Sam, K. (2011, June 6). *Storage Configuration*. Retrieved March 11, 2013, from http://trac.zentyal.org/wiki/Documentation/Community/HowTo/ThePerfectEboxSetup/StorageConfiguration

Sammons, J. (2012). *The basics of digital forensics: The primer for getting started in digital forensics*. Waltham, MA: Syngress.

SANS (2013, June 28). *SANS Institute: About*. Retrieved June 28, 2013, from http://www.sans.org/about

Sawyer, J. J. (2006). MAGNETIC DATA RECOVERY - THE HIDDEN THREAT. doi:10.1.1.87.3037

Schmid, M. (2009, March 2). How an STM Works. Retrieved from http://www.iap.tuwien.ac.at/www/_media/surface/stm_gallery/stm_animated.gif

Seagate (2011). *The Top 20 Things to Know About SSD* (ssd-faq-tp612-2-1103us). Retrieved from Seagate Technology LLC website: http://www.seagate.com/files/www-content/product-content/pulsar-fam/_cross-product/en-us/docs/ssd-faq-tp612-2-1103us.pdf

SEI (2010, April 21). *Trademarks and Service Marks*. Retrieved August 4, 2013, from http://www.sei.cmu.edu/legal/marks/

Shimpi, A. L. (2009, December 31). *AnandTech | OCZ's Vertex 2 Pro Preview: The Fastest MLC SSD We've Ever Tested*. Retrieved April 6, 2013, from http://www.anandtech.com/show/2899

Sood, A., James, G. M., Tellis, G., & Zhu, J. (2012). Predicting the Path of Technological Innovation: SAW vs. Moore, Bass, Gompertz, and Kryder. *Marketing Science*, *31*(6), (pp. 964-979). doi:10.1287/mksc.1120.0739

Tal, A. (2002). *NAND vs. NOR flash technology* (02/01/2002). Retrieved from Electronic Products website: http://www.electronicproducts.com/Digital_ICs/NAND_vs_NOR_flash_technology.aspx

Tech Juice (2011). *An Introduction to Hard Disk Geometry | Tech Juice*. Retrieved from http://www.tech-juice.org/2011/08/08/an-introduction-to-hard-disk-geometry/

Travis, J., & Rau, R. M. (2004). *NIJ Special Report, Forensic Examination of Digital Evidence: A Guide for Law Enforcement, April 2004* (NCJ 178280). S.l: U.S. Department of Justice Office of Justice Programs.

Wei, M. Y., Grupp, L. M., Spada, F. E., & Swanson, S. (2011). Reliably Erasing Data from Flash-Based Solid State Drives. *In Proceedings of the 9th USENIX conference on File and stroage technologies (FAST'11)*, (pp. 8-20).

Wilcoxon (1994). *Bearing failure: Causes and cures*. Retrieved from Barden Corporation website: www.wilcoxon.com/product_presentations/bearing.pdf

Wilkinson, S. (2012). *Good Practice Guide for Computer-Based Electronic Evidence*. Retrieved from 7safe website: http://www.7safe.com/electronic_evidence/ACPO_guidelines_computer_evidence.pdf

Williams, J. (2007). The ACPO Good Practice Guide for Managers of e-Crime investigation. Retrieved from http://www.acpo.police.uk/documents/crime/2011/201103CRIECI14.pdf

Winter, R. (2013). SSD vs HDD – data recovery and destruction. *Network Security*, *2013*(3), (pp. 12-14). doi:10.1016/S1353-4858(13)70041-2

Wright, C., Kleiman, D., & Sundhar, S. (2008). Overwriting Hard Drive Data: The Great Wiping Controversy. *ICISS 2008*, *5352*(0302-9743), 243-257. doi:10.1007/978-3-540-89862-7_21

Wu, J. (2003, November 18). Electrical adapter for connecting connectors of different interface. Retrieved from http://www.freepatentsonline.com/6648695.html