ADISA
ASSET DISPOSAL & INFORMATION SECURITY ALLIANCE

White Paper

University of South Wales
Prifysgol De Cymru

Managing End of Life Risk

For Solid State Devices (SSD).

Professor Andrew Blyth and Steve Mellings

November 2014
V1.5

Abstract

Solid state devices (SSD) have revolutionised user technology and the use of products utilising this storage is now widespread within most businesses. However, when those devices that use SSD reach their end of life there is a range of challenges when looking to meet data protection requirements. The greatest challenge is that there is currently no government approved overwriting solutions for data sanitisation on devices using SSD which leads many end users to ask; what do I do to protect my data when releasing these assets at end of life?

Authors

Professor Andrew Blyth, PhD.
Information Security Research Group,
University of South Wales, Treforest, RCT, CF37 1DL
andrew.blyth@southwales.ac.uk

Steve Mellings,
Founder,
ADISA,
Hamilton House, 1 Temple Avenue, London. EC4Y 0HA
steve.mellings@adisa.org.uk

Introduction.

It was Sir Francis Bacon who said, "*Knowledge is power*", and in today's technology driven world we can map this into information and through information into data. The critical nature of this relationship has been recognised via a number of national governments, manifesting itself in a number of directives, pieces of law and strategic statements all relating to data management, information security or privacy.

Within the United Kingdom, Objective 3 of the UK National Cyber Security Strategy 2011 [1] is:

> "*Helping to shape an open, vibrant and stable cyberspace, which the UK public can use safely and that supports open societies.*"

And so we can see that there is a battle between the desire for a vibrant and open cyberspace, maximising the potential of information sharing and a legal requirement for businesses to protect data pertaining to the individual, or to keep their own corporate information private. Within UK Law it is the Data Protection Act 1998 that legislates this area. This is governed by a set of principles relating to data protection with Principle 7 being the relevant criteria relating to how data, and data bearing assets, should be managed at end of life [2]:

> *"Principle 7: Appropriate technical and organisational measures shall be taken against unauthorised or unlawful processing of personal data and against accidental loss or destruction of, or damage to, personal data."*

The concept of 'Appropriate' behaviour is a repeating one in both industry and other country specific law and deeming what is appropriate is perhaps the key question. Within the heart of this debate are operations put in place, aimed to protect information and data. These operations, including the technology itself, should be specified, deployed, managed and disposed of appropriately in order to meet regulatory requirements.

Within the past ten years the challenge of displaying 'appropriate' behaviour has grown exponentially, as how and where we work has changed dramatically. Users now demand the ability to work where they like, which has seen technology move from desk based devices to the multitude of different platforms and devices now being used in the work place. Manufacturers have re-energised the market by introducing smaller, lighter, but more powerful devices, which has allowed more productive working practices to evolve. Consequently working on trains, in hotels and at home have now all become part and parcel of everyday life, which has forced information security and within that, data protection, to become a multi-layered, multi-faceted and fluid process.

A key component in facilitating this change in working practice has been the evolution of NAND based storage technology. Commonly referred to as Solid State Drives (SSD) [3], this small form factor storage has allowed devices to become more portable, faster and more utilitarian. Adoption of this technology has led to increased user benefits and productivity, but has also led to a series of challenges later within the product lifecycle, which were perhaps not considered at the point of deployment. Issues surrounding device security and management are well documented but new concerns have emerged regarding the retirement of these assets with the primary concern being simple, how to sanitise the data?

At this moment in time there are no government approved overwriting products for this media type, which is leaving businesses to make a decision on their own in regard to how best to release assets utilising this media from their estate. This lack of guidance is forcing a position of either uncontrolled risk taking or the adoption of a risk avoidance approach at end of life leading to the destruction of perfectly serviceable devices.

Growing legislation and regulatory requirements.

Within this operational theatre, corporate governance is becoming a far more onerous requirement and the need to show compliance requires more than a simple tick box mentality. When we consider 'privacy' or 'data protection' there has been a growing shift to not only legislate, but also to bring to bear regulatory actions. In Europe the EU Data Protection Directive 95/46/EC is currently undergoing a complete rewrite aimed at introducing an EU wide law, which will help bring legislation right up to date in regard to the way in which technology and data is used and managed today. A key part within the business world is that the regulators will now be able to wield a much larger stick should those companies who control or process personal data choose to ignore their responsibilities, as fines of up to 2–5% of global turnover have been proposed. Within the United States, whilst there is no national authority for data protection, there are over twenty sector specific privacy or data security laws and hundreds of state laws (California has more than twenty-five state privacy and data security laws itself [9]). Industry specific regulators and laws such as the Federal Trade Commission (FTC) and the Health Insurance Portability and Accountability Act of 1996 (HIPAA), are perhaps the most active in the area of privacy and data protection and they also offer a more prescriptive direction over acceptable organisational and technical controls to be followed.

So for those individuals who have data protection, information security, compliance, governance or even brand protection within their remit, all areas where their own company could be compromised needs to be reviewed. We can see that where data stored on technology is concerned, the process of end of life asset disposal MUST be included with the key objective of that process being data sanitization. Within this area, thanks to the changes in technology an understanding of SSD, now used so widely on smart phones, tablets and laptops, will be crucial to writing policy and service specifications that identify and mitigate the risk this media poses at end of life.

This paper is written as an aid for general IT experts to help understand the challenge of securely erasing SSD at end of life. It explores the technology with specific reference to those devices that use the ATA command set function. It then expands on the challenge of overwriting these devices, how to validate a successful overwrite and finally gives businesses clear guidance on what to do in order to meet their own compliance needs and brand protection desires. In order to ensure the audience are familiar with some concepts, the paper also introduces and outlines risk management including threat profiling.

Introducing the concept of risk management.

There are bold statements such as “We don’t allow risk” or “Our attitude to risk is that there is none”, but the reality is that risk is an everyday part of normal business practice. It is how this risk is managed which is critical and on a personal level, whilst we may not know it, each and every one of us performs risk assessments almost every day. A simple act such a crossing the road where there is no official crossing point is an exercise in risk management.

For businesses the notion of understanding operational risk and putting in place processes and countermeasures to manage that risk is second nature. This is not restricted to those involved in information security, but extends into many other areas, such as facilities management, human resource management and logistics. Whether as a result of corporate need or regulatory necessity, risk management is a process businesses have to embrace and control.

The term risk is defined as follows from ISO 13335/1/2004.

> *“A risk is a potential that a given threat will exploit vulnerabilities of an asset or group of assets and thereby cause harm to the organization. It is measured in terms of a combination of the probability of an event and its consequence.”*

Within asset retirement, there is a whole array of internal and external risks that need to be identified, assessed and managed accordingly, such that the objective of the process is achieved. A key part of any risk management is the notion of threat. Understanding where threat may come from and what motivation and capability that threat actor may have are essential when looking to build risk mitigation processes. Without this understanding the type and location of countermeasures is impossible to define, which means processes may be insufficient or in other cases too significant, causing needless cost or loss of opportunity.

Understanding Threat.

For the purposes of this paper we will focus on just one risk, which is the risk of data being accessed after an overwriting technique has been deployed on that media. As such, in order to mitigate that risk we must first review the threat that could cause that risk to become a reality. In order to help assess the threat of data recovery, ADISA and the University of South Wales have standardised a threat matrix, which outlines capabilities of a particular threat adversary so that sanitisation techniques can be measured against that threat actor.

The Threat Matrix.

The threat matrix, on page six, defines a series of capabilities and risks that various threat actors can pose against the recovery of data from any media.

Risk Level	Threat Actor and Compromise Methods
1 (Very Low)	Casual or opportunistic threat actor only able to mount high-level non-invasive and non-destructive software attacks utilising freeware, OS tools and COTS products.
2 (Low)	Commercial data recovery organisation able to mount non-invasive and non-destructive software attacks and hardware attacks.
3 (Medium)	Commercial computer forensics organisation able to mount both non-invasive/non-destructive and invasive/non-destructive software and hardware attack, utilising COTS products.
4 (High)	Commercial data recovery and computer forensics organisation able to mount both non-invasive/non-destructive and invasive/non-destructive software and hardware attack, utilising both COTS and bespoke utilities.
5 (Very High)	Government-sponsored organisations or an organisation with unlimited resources and unlimited time capable of using advanced techniques to mount all types of software and hardware attacks to recover sanitised data.

Table 1 – The Threat Matrix

The use of this threat matrix allows both industry and end users to look at their own sanitisation choice and assess whether they are fit for purpose. For the act of sanitisation each end user/data controller should assess where they believe their threat will come from and what capabilities they may have. A key part of this assessment will be the value of the data itself. If the data being sanitized would have value to someone else, then they will be viewed as your threat actor. So if your data has little value, other than perhaps for embarrassment via press, then perhaps your threat actors will be a risk level 1. If your data is more valuable and there are entities that seek it in a positive fashion, then they may mount more pro-active attacks using different capabilities. The type of attack will directly correlate to the value of the data, so companies holding government data would be viewed at the higher level of risk, as this type of data is sought by competing governments, the press and perhaps large corporates.

For those who provide data sanitisation services the tendency is to assume higher level attacks, so that they are protected should their customer's data be of that level. This can often lead to over protectionism and loss of opportunity to re-use assets. This is precisely where we are today when looking at solid state media, as with no government approved overwriting tools, both the industry and end user community look at each other for the answers when the question of SSD re-use is raised with neither party confident of making the required decision.

Now that we have outlined the current position, let us move the discussion forward so we understand why there are no government approved solutions for SSD and present steps that can be taken to help provide some of those required answers.

The Solid State Storage Device Technical Architecture.

To understand the issues let us first introduce the technology itself. Solid state technology can be utilised using ATA or SCSI command sets in solid state hard drives or smart phones and tablet technology using a USB interface. The majority of SSD devices seen in the field utilise the ATA command set.

The conceptual architecture for a SSD utilising ATA command sets [3] is defined in Figure 1 and comprises three components, The Host Interface Logic, The SSD Controller and the NAND Storage. The roles of each of these may vary depending on the age of the SSD.

The **Host interface Logic** defines how a SSD will function as a computer hard-drive [1] via an ATA interface. The role and function of the Host Logic Interface is to define the interface between the device and the ATA / SCSI command unit bus and thus to make the NAND storage unit appear as a local block addressable (LBA) device to the user on the ATA / SCSI interface/bus.

The **SSD Controller** manages the Flash Translation Layer (FTL) and Data Compression of data being written to the **NAND storage**. The SSD controller comprises:

- The processor, which is responsible for managing the system bus and receiving and processing commands from the Host Interface logic. The processor is also responsible for managing the AES crypto keys.
- The AES Crypto Device is responsible for performing the AES-256 bit crypto function.
- The Buffer Management provides a buffer for data being read/written to the data bus with the SSD controller.
- The Flash Controller implements the Flash Translation Layers.

To manage the storage located on the NAND chip, the NAND storage devices function to provide a data storage capability that can be read/written onto the NAND data bus.

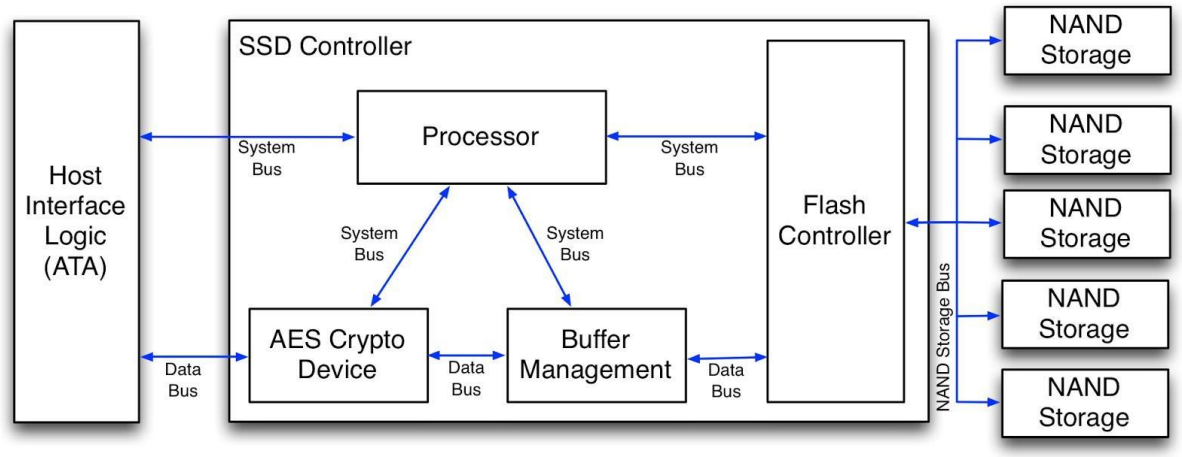

Figure 1 – SSD Architecture

Depending on the make and model of the SSD, the following functions/operations are utilised:

- During the write operation data will be stripped across the different flash chips in the device;
- Compression algorithms will be employed;
- Duplication algorithms will be employed.

The SSD controller makes use of technologies such as the Flash Translation Layers (FTL), Data Compression and AES encryption to ensure that the data is write/read from the NAND/Flash storage chips. The FTL functions to make linear flash memory appear to the system like a disk drive. It does that by doing a number of things [3]. First, it creates "virtual" small blocks of data, or sectors, out of flash's large erase blocks. Next, it manages data on the flash so that it appears to be "write in place", when in fact it is being stored in different spots in the flash. The data compression functions to ensure that the maximum amount read-writes on the device can be achieved, and the AES encryption is used to ensure that all data written to a NAND storage cell is encrypted, and all data read to a NAND storage cell is decrypted. AES encryption makes use of a 256-bit key. The data bus functions to allow the controller chip to read and write data to/from the storage devices located on the NAND storage bus [3].

Over provisioning and wear levelling.

In order to extend the life of SSD, manufacturers build the media with a greater capacity of storage than the stated amount. This allows the controller chip to spread the volume of read-writes across the wider capacity so that the overall life of the device is longer. This process is typically called wear levelling and is managed by the execution of an algorithm, which controls where the data is mapped. There are two basic types of wear levelling mechanisms used in flash memory storage devices:

- Dynamic wear levelling.

  Dynamic wear levelling uses a map to link logical block addresses (LBA) from the ATA interface to the physical Flash/NAND storage. Each time the operating system performs an ATA write command, the map is updated so the original data block is marked as invalid data, and the new block is linked to that map entry. Each time a block of data is rewritten to the Flash/Nand memory, it is written to a new location. The drive will last longer than one with no wear levelling, but there are blocks still remaining as active that will go as unused when the drive no longer functions. (However, blocks that never get replacement data sit with no additional wear on the Flash memory. The name comes from the fact that only dynamic data is being recycled. The drive may last longer than one with no wear levelling, but there are blocks still remaining as active, that will go unused when the drive is no longer operable. Typically when writing data to a NAND cell a SSD will make use of compression to optimize the utilisation of the NAND storage and minimize the number of access read/writes to a cell.)

- Static wear levelling.

  Static wear levelling also uses a map to link the LBA to physical memory addresses. Static wear levelling works the same as dynamic wear levelling except the static blocks that do not change are periodically moved so that these low usage cells are able to be used by other data. This rotational effect enables an SSD to continue to operate until most of the blocks are near their end of life. This is also sometimes referred to as global wear levelling, as the entire image is levelled.

Data Sanitisation Methods.

When looking to sanitise data on a range of media types there are a number of processes available. The choice of which one to choose will depend on the risk assessment, the media type and of course, cost.

Software Methods: The Traditional Method.

The traditional software method for data sanitisation is to overwrite the data using a known string of data, typically 0xFF and/or 0x00. Over the years various standards have been developed that say that data should be written to every LBA on a HDD as well as the HPA and DCO between 1 and 7 times. The NIST special publication 800-88 on Guidelines for Data Sanitization [8], outlines the traditional approach to data overwriting and for "Clear" recommends that the data is to be written multiple times to a device so as to ensure that all addressable areas have been overwritten. The implementation of this approach to data overwriting makes use of traditional ATA commands supported by every ATA capable device. In effect all traditional methods make use of ATA commands to READ_LBA and WRITE_LBA [3]. These ATA commands are used all the time by the operating system, as they are the basic structures through which data is read and written to a device.

Software Methods: Secure Erase.

As each ATA device implements a set of ATA commands, some ATA devices (such as certain SSDs) implement the ATA secure erase function [3]. This feature can be invoked at the command line using tools such as hdparm. The following is an example of how we can use hdparm to issue the ATA secure erase command to a device/dev/ssd1, give the master password of P55w0rd on the drive.

```
hdparm –user-master u –security-erase Pa55w0rd /dev/ssd1
```

ATA Secure Erase is part of the ATA ANSI specification and when implemented correctly, wipes the entire contents of a drive at the hardware level, instead of through software tools. The current ATA specification for Normal Erase mode states that the SECURITY ERASE UNIT command shall write binary zeroes to all user data areas.

Software Methods: Encryption.

An increasing number of solid state drives are making use of AES 256-bit encryption inside the hardware controller. This means that before data is written to a NAND cell, the data is first encrypted. The encryption key is stored in an obfuscated manner on the controller chip. This approach to data protection ensures that hardware attacks to recover data such as chip-offs will fail, as all they will be able to recover from the NAND cell is encrypted data.

The advantage of such an approach to data protection means that data can effectively be rendered irrecoverable by the controller chip resetting the AES encryption key to a new value. Technically the data is still present in the NAND cells; it is just that without the key to decrypt it, to all intents and purposes the data is unrecoverable. [3] As long, of course, that the encryption algorithm remains intact.

Hardware Methods: Physical Destruction.

The objective of physical destruction is to badly warp, distort or destroy the device, rendering the drive or any component of the drive inoperable. SSD require a different approach than magnetic hard drives, as the data storage area are the NAND cells, are far smaller and have no moving parts, so some of the more basic physical

approaches such as the use of hammers will not be totally effective. Any physical means of sanitisation should ensure that all NAND cells are physically attacked, so that any hardware efforts to effect recovery are rendered impossible.

Hardware Methods: Degaussing.

Degaussing is the process of decreasing or eliminating a remnant magnetic field. It is possibly named after the gauss unit of magnetism. Due to Magnetic Hysteresis it is generally not possible to reduce a magnetic field completely to zero, so degaussing typically induces a very small “known” field, referred to as bias. Degaussing is used to reduce magnetic fields in CRT monitors and to destroy data held on magnetic data storage, but for SSDs the theoretical field strength required far exceeds current commercial products rendering degaussing ineffective.

One of the challenges with degaussing is that the strength of the field determines the effectiveness of the degaussing operation on the device, meaning correctly calibrated equipment is essential. Also an appreciation of the physical nature of the device is essential, so that any material that could act as a shield to the magnetic field is taken into considering. A further challenge is that there is no physical change in the media once degaussed, leaving operators reliant on process control and/or random quality checks to confirm that a device has been degaussed. Finally, the magnetic field is indiscriminate in the effects that it has and any device touched by the magnetic field may be affected by that field.

Hardware Methods: Aggregation.

Whilst not a means of sanitisation itself; aggregation is a useful aid when considering the decrease of risk within end of life processing. Aggregation is the process of introducing a volume countermeasure when defending an asset from an attack. Not viewed as a singular form of defence, aggregation can help protect a known number of assets by introducing that range into a larger group of assets, thus making the sample require for a successful attack much greater. This is of particular importance when looking at creating a defence against a higher level of attack. If the asset which could be subject to a forensic level of attack is ‘hidden’ amongst a much larger sample of similar assets, then the resource required by a threat adversary to identify and then attach that asset is much higher. This may make this type of attack less viable than a more direct attack.

What are the challenges of data overwriting?

SSD are storage devices, which utilise NAND cells for storage and controller chips for device management and user interface. (See Figure 1) A key part of SSD architecture is over provisioning, which is where the total storage within each device is greater than the available storage to the user and is intended to extend the life of the device. There are a range of technical functions that happen during the operation of SSD including wear levelling, garbage collection and data compression. During in-life use these process are at the very core of the benefit of using SSD, but at end of life they become a significant handicap when validating traditional data overwriting techniques. Traditional overwriting (commonly referred to as data erasure) is achieved by writing a series of characters to all addressable areas of a magnetic hard drive. Validating that this has occurred is easy as a sample of sectors can be read and confirmation that either (a) there is no information present, or (b) that consistent overwriting patterns are in place, this is not so easy when referring to SSD. For SSD overwriting there are different approaches, three of the most common are:

- A known pattern of data is written from the start of the device to the end of the device. The number of times that this process is repeated is derived from the implementation of the wear levelling algorithm. The algorithm is implemented differently by each vendor and on different devices from the same vendor, so the number of necessary overwrites to successfully remove all data varies.
- A SSD USB controller chip implements a secure erasure command, which will then either make use of a software solution or a hardware solution to erase all of the data on the NAND storage cells.
- A form of crypto erase is possible, whereby the encryption key is located and erased. This is an extremely challenging approach as for obvious reasons the location of these keys is not widely known!

The key problem when dealing with SSD is how do we validate that the overwriting process has been successful? A successful overwrite could be measured in two ways; (a) that all physical memory locations have been included in the overwrite and (b) that there is no data recoverable using forensic techniques. Confirming a successful overwrite is an issue, as there is no one-to-one mapping of data locations due to wear levelling. Furthermore when a cell fails or reaches its maximum read/write it becomes degraded and not addressable. This effectively leaves data still resident on such cells after an overwrite has been instigated through the controller chip, thus rendering the term “all data areas” impossible to comply with.

For this reason existing approval schemes, such as CPA run by CESG, cannot validate an overwriting software as arguably there could be degraded cells on a device and as such data may remain. A further issue is that SSD is a new technology and the manufacturers of these devices do not follow consistent protocols during manufacturer. Chip substitution is commonplace, which means that the same brand of SSD may have different manufacturers’ components inside. As the controller chip functions are potentially implemented differently, depending on manufacturer, then the question of “what SSD do I have?” must become one of “what chips sets are within my SSD estate?” As such it is essential to consider the issue of sanitisation at point of procurement to ensure that the SSD product sets within the estate are a common build to make end of life processing easier.

So to summarise; the challenge of overwriting is how to address over provisioning, how to handle degraded cells and finally, how can we validate that overwriting has successfully been implemented.

The Solution When Dealing with End of Life SSD?

Confused? At this stage most companies when they are aware of these issues will say, “Destroy”. It’s the easiest process to implement and allows companies to manage their risk. However, to simply destroy assets, which in some cases, will hold up to 25% of the original value, is wasteful, not only in an environmental sense but also financial. So what to do?

Unlike overwriting for magnetic hard drives, there is no single absolute statement, which can be relied on to make you, the risk owner, feel comfortable. As such each risk owner must make an informed, risk based decision on how to deal with SSD at end of life. This may sound daunting but the good news is that by following the steps below, we believe that by adopting a risk base approved to SSD overwriting, that you will be able to understand the risk and build in the correct procedural and technical countermeasures.

There are **5 key stages to secure asset disposal** and these should be followed regardless of the media type. These are:

- Stage 1: Policy Development.
    - Display management control through a prescriptive asset disposal policy. This should include data categorization, business impact, threat profiling, risk assessment and finally, this should produce an approved media sanitisation profile. (See Appendix B.)
- Stage 2: Organisational Control.
    - Process and Procedure documents are required to help deliver policy. For this reason, any media that is being disposed of should be done so in a controlled way. A phrase used by regulators both in the UK and the US is “Organisational control”. Clearly if there is no control over the process, then can any company confidently state that they can show “organisational control?”
- Stage 3: Third Party Control.
    - Control all external engagements with a clearly defined service specification as part of a written contract. This should include ALL potential outputs from the business, including end of life, end of lease and mid-life instances such as repair. Look for relevant standards, which show independent assessment of their security capabilities. [11]
- Stage 4: Compliance.
    - To be able to show compliance with your own policy (which in turn should show corporate compliance to regulatory requirements), a thorough audit programme is required which should result in a reporting schedule able to evidence control over this process.
- Stage 5: Review and Reflection.
    - Technology changes, partners change and threats change, so it is essential to review and consider approaches to asset disposal at regular intervals.

As we are talking specifically about SSD let us drill down into Stage 1 in greater detail and in particular how to decide on an approved means of sanitisation for SSD.

Step 1 to building an approved means of sanitisation for SSD.

- Data Categorisation:
    - It is essential to consider the category of data, which your company controls. At a superficial level if it is data that sits entirely in the public domain, then the category would be very low (it is publicly available anyway). However, most businesses hold data pertaining to an individual (pay roll for example) or their own corporate data, which ensures that their data should be

handled in a protective manner. How data is categorised is entirely up to the company and there will be different approaches to data protection depending on the data category. Perhaps a three-phased approach of public, official or secret could be used?

- Business Impact:
    - Once categorized, the impact to the business of the breach of that data should also be assessed. Whilst the payroll of the CEO would be categorized as highly sensitive, its breach would only cause minor impact to the business (perhaps with some shareholder dissonance!) whereas the loss of the entire staff payroll would cause significant impact to the business. Almost certainly resulting in regulatory action, the fines associated with such a breach would be significant, but also the cost of addressing the breach would be equally significant.

- Threat Profiling:
    - Who potentially could be targeting you? Thanks to the proliferation of the dark net, hacking skills (SKS Cyber skills) are now being shared, meaning that the volume of bedroom hackers has grown exponentially, which has largely been the cause in the proliferation of cyber-attacks being carried out. Organised crime has also been quick to skill up in technology and are ever more resourceful when looking to seek new ways of sourcing funds. So who potentially is targeting you? What is their motivation? What are their skills?

Step 2: Identification of SSD assets.

It is essential to identify where SSD sit within your business, such that all possible outputs are aligned to the same policy. Product sets such as laptops, tablets and smart phones need to be included as well as potential hidden SSD within networking equipment or even printer technology. Regardless of the media being discussed, the creation and management of the chain of custody IS ESSENTIAL and without this being in place not only is asset leakage inevitable but regulatory compliance impossible to prove.

Step 3: Risk Assessment.

Once we have gone through these stages then a risk assessment can be made. This assessment should look at the processes surround asset retirement as well as the endpoint sanitisation technique required. As we are focussed just on the act of sanitisation we should re-introduce your options.

Physical Destruction; some traditional hard drive destruction techniques won't work on SSD simply because of the physical nature of the product. Some larger shred sizes may damage the boards but could miss the NAND cells themselves resulting in potential cell level recovery. SSD destruction should be achieved in such a way where every NAND cell is impacted and for higher levels of attack, the potential requirement for them to be disintegrated may come into play.

Degaussing; theoretically if a magnetic waveform was so strong and used on a SSD it could erase SSD as it could force all the electrons within the NAND cells to revert to a single state. However, current commercial degaussers do not work at this high level (and it is unlikely they ever will) and therefore won't work on SSD.

Overwriting; in 2013 we published a methodology for testing over-writing solutions used on SSD [10]. During the tests we have carried out by key vendors of SSD overwriting tools we have discovered that based on various degrees of forensic attacks, data cannot be recovered from a particular set of SSD after being overwritten. It is worthwhile reviewing this methodology and looking at some of the results, which have been found on commercially available products. Sadly we have been unable to offer confidence that "Product X works on all SSD", because we have found SSDs (particularly earlier SSDs) to be manufactured differently and also that our

assurance of the inability to perform data recovery is based not on conclusive evidence of an overwriting pattern, but that at a cell level hardware encryption was being invoked, which meant that even after intrusive and destructive attacks, all that could be seen was encrypted data. (Whether than is the original data pattern or overwrite is impossible to gauge.)

Step 4: Decision on the approved means of sanitisation.

At this stage there should be an understanding of the category of the data, the impact of that data loss and where the threat may be coming from. A risk assessment can now be made for deciding on the appropriate means of sanitisation. For those seeking to promote the re-use of assets we would always recommend that any decision should include the use of products which have been forensically tested. This forensic testing can follow a similar method which ADISA advocates or can follow a method which the risk owner themselves approves. However, due to the nature of the technology it is essential than testing is carried out. For legal compliance it is also recommended to seek products which offer the additional benefit of an audit trail (serial number reporting) and indemnity insurance.

An example of an approved means of sanitisation is shown in Appendix B and it is crucial to look at the processes included BEFORE sanitisation. Like all security, the biggest vulnerability lies not in the technology, but in the people around it, so don't forget that inventory management (chain of custody), partner management and verification of the service being executed is perhaps the most important part of all. After all, if a device doesn't make it to the bench to be sanitized, then all of the above is academic!

Closing Comment.

Most organisations when asked are not aware of the issues of end of life SSD processing and so the issue only becomes critical at the time when the process is required. This short time line results in risk avoidance rather than risk management and so the key to SSD at end of life is to engage in the process BEFORE you need it!

Data controllers should engage with recognised software developers, engage with their IT asset disposal (ITAD partners) and engage with forensic experts. **Together**, you CAN control risk and ensure that you meet the requirements of the data / privacy regulators, your own data protection requirements, whilst at the same time promoting re-use and benefiting from the residual value lock within these assets.

Whilst content of this paper is applicable to any flash based storage device, the conclusions should only be considered with reference to devices that are utilising the SSD controller chips to access Flash/NAND storage cells. For products such as smart phones we would reference research undertake by the author, which will be released in March 2015. Devices such as USB sticks, digital cameras and smart phones that do not utilise SSD controller chips to access NAND storage cells many of the research findings in this paper will not be applicable.

References

[1] The UK Cyber Security Strategy, Protecting and promoting the UK in a digital world, UK Crown Copyright, November 2011.

[2] The UK Data Protection Act, 1998. http://www.legislation.gov.uk/ukpga/1998/29

[3] Rino Micheloni, Alessia Marelli and Kam Eshghi, “Inside Solid State Drives (SSDs)”, Springer Series in Advanced Microelectronics, Springer, 2012.

[4] Open NAND Flash Interface Specification, Version 3.0, March 2011.

[5] Open NAND Flash Interface Specification: Block Abstracted NAND, Version 1.1, July 2009.

[6] ATA Attachment-8 – Serial Transport (ATA8-AST), Aug 2011.

[7] Understanding the Flash Translation Layer (FTL) Specification and Data Compression, Intel Application Note, AP-684. 1998.

[8] Richard Kissel, Matthew Scholl Steven Skolochenko, Xing Li, NIST Special Publication 800-88: Guidelines for Media Sanitization, NIST, 2006.

[9] Data Protection Laws of the World. DLA Piper. April 2014.

[10] ADISA Solid State Testing Methodology v 1.0 (June 2012) and ADISA Claims Testing Methodology v 1.0 (January 2014). Available from www.adisa.org.uk/claimstesting/

[11] ADISA IT Asset Disposal Industry Security Standard 2013 v1.4. Available from www.adisa.org.uk

Appendix A – ADISA Risk Methodology

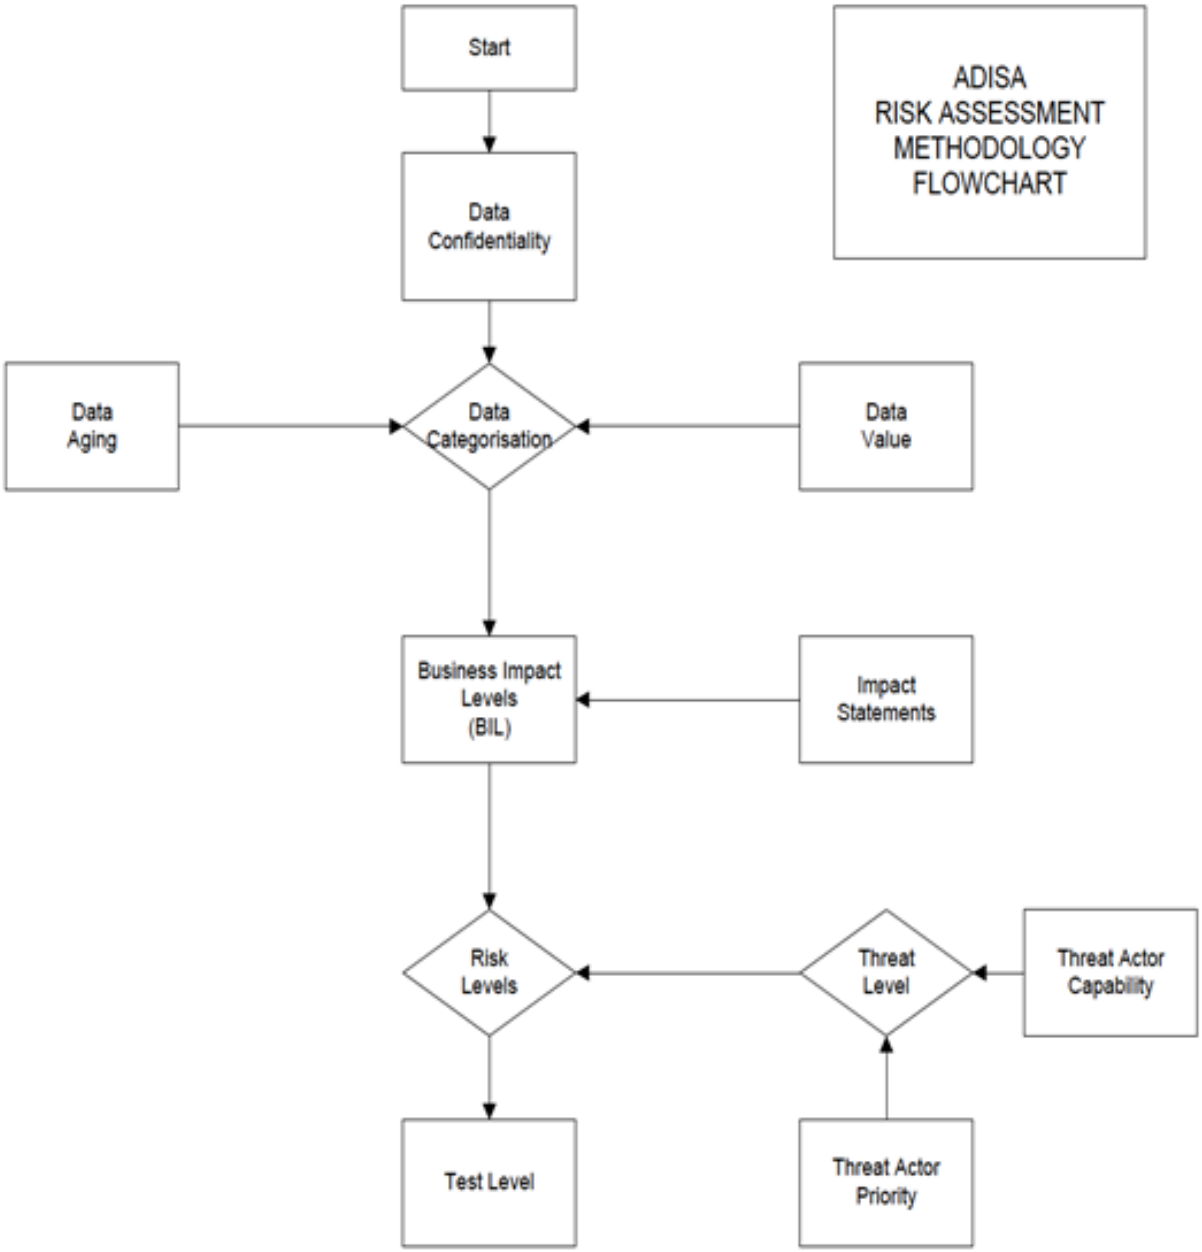

Appendix B – ADISA example of an approved means of sanitisation

NB: This is an example and is purely for illustrative purposes.

Media Type	Product Set	Risk Level	At Office Site	At Service Provider	Re-Use?
Solid State	Smart Phone	5 to 6	Inventory phone using IMEI number, invoke manufacturers reset, book into secure stores, book out of stores, destroy on-site using mobile shredder to 22mm.	n/a	No
Solid State	Smart Phone	1 to 4	Inventory phone using IMEI number, invoke manufacturers reset, book into secure stores	Validate receipt against inventory list. Utilise software that has undergone testing and offers indemnity and audit capability.	Yes
Solid State	Tablet – Apple	All	Inventory device using serial number, invoke manufacturers reset, book into stores.	Validate receipt against inventory list. Utilise software that has undergone testing and offers indemnity and audit capability.	Yes
Solid State	Tablet – Other	All	Inventory device using serial number, book into stores.	Validate receipt against inventory list. Utilise software, that has undergone testing and offers indemnity and audit capability.	Yes
Solid State	Laptop	All	Inventory device using serial number, book into stores.	Validate receipt against inventory list. Utilise software that has undergone testing and offers indemnity and audit capability.	Yes

Appendix C – Current results of testing undertaken against ADISA SSD testing methodology

Reference	Vendor	Product Details	Test Product	Test Level
ADPC0001	Blancco	Blancco 4.x	Micron RealSSD 128Gb	1
ADPC0002	Blancco	Blancco 4.x	Crucial M4 128GB	1
ADPC0003	Blancco	Blancco 4.x	Intel 80GB	1
ADPC0004	Blancco	Blancco 4.x	Kingston 128Gb	1
ADPC0006	Blancco	Blancco 4.x	Samsung 128Gb	1
ADPC0008 3.1	Tabernus	Enterprise Erasure v5.3.20	Micron RealSSD C400128Gb	1
ADPC0008 3.2	Tabernus	Enterprise Erasure v5.3.20	Intel SSD 520 120Gb	1
ADPC0008 3.3	Tabernus	Enterprise Erasure v5.3.20	Micron RealSSD P400m 2.5 100Gb	1
ADPC0008 3.4	Tabernus	Enterprise Erasure v5.3.20	Seagate 120Gb	1
ADPC0008 3.5	Tabernus	Enterprise Erasure v5.3.20	Intel X25-V 40Gb	1
ADPC0008 3.6	Tabernus	Enterprise Erasure v5.3.20	Toshiba 128Gb	1
ADPC0011	Tabernus	Enterprise Erasure v7.1	Toshiba 128Gb SSD	1
ADPC0012 1.1	ITRenew Inc.	Teraware v2.15	HP 200Gb SAS-SSD	1 and 2
ADPC0012 1.2	ITRenew Inc.	Teraware v2.15	HP 400Gb SAS-SSD	1 and 2
ADPC0012 1.3	ITRenew Inc.	Teraware v2.15	Intel 320 Series 160Gb SATA-SSD	1
ADPC0012 1.4	ITRenew Inc.	Teraware v2.15	(Dell) Intel 320 Series 160Gb SATA-SSD	1
ADPC0012 1.5	ITRenew Inc.	Teraware v2.15	Intel X25-M 160Gb SATA-SSD	1
ADPC0012 1.6	ITRenew Inc.	Teraware v2.15	Samsung PM800 Series 128Gb SATA-SSD	1
ADPC0012 1.7	ITRenew Inc.	Teraware v2.15	Samsung PM810 Series 128Gb SATA-SSD	1
ADPC0012 1.8	ITRenew Inc.	Teraware v2.15	Samsung SS410 Series 32Gb SATA-SSD	1

Appendix D – ADISA Threat Matrix and Test Levels

The Threat Matrix.

The threat matrix defines a series of capabilities and risks that various threat agents can pose against an asset. The test levels define a series of capabilities that a threat actor/agent may wish to bring against an asset either by direct access to the asset or access via its location within a device.

ADISA Risk Level	Threat Actor and Compromise Methods	Type of Method	ADISA Test Level
1 (Very Low)	Casual or opportunistic threat actor only able to mount high-level non-invasive and non-destructive software attacks utilising freeware, and OS tools.	Keyboard attacks from a motivated individual. Typical attack could be using open-source forensic tools.	1
2 (Low)	Commercial data recovery organisation able to mount ADISA Risk Level 1 attacks and non-invasive and non-destructive COTS software attacks and hardware attacks.	Keyboard attacks from a motivated professional organisation. Typical attack could be using commercial tools.	1
3 (Medium)	Commercial computer forensics organisation able to mount ADISA Risk Level 2 attacks and invasive/non-destructive software and hardware attack, utilising COTS products.	Laboratory attacks from commercial data recovery experts. Typical attack could be: Chip Readers/bus decoders.	2
4 (High)	Commercial data recovery and computer forensics organisation able to mount ADISA Risk Level 3 attacks and invasive/destructive software and hardware attack, utilising both COTS and bespoke utilities.	Laboratory attacks from specialist forensic scientists. Typical attack would involve analysis of individual hardware components.	2
5 (Very High)	Government-sponsored organisations using advanced techniques to mount all types of software and hardware attacks with unlimited time and resources to recover sanitised data.	An attack agent of unknown capability and unlimited resource. Typical attacks: Taking theoretical forensic possibilities and making them an actual capability.	3

Table 1 – The Threat Matrix