




May 18th, 11:15 AM - 11:50 AM

Analysis of Data Erasure Capability on SSHD Drives for Data Recovery

Andrew Blyth
Technology Research Centre, DPG

Follow this and additional works at: <https://commons.erau.edu/adfsl>

 Part of the [Aviation Safety and Security Commons](#), [Computer Law Commons](#), [Defense and Security Studies Commons](#), [Forensic Science and Technology Commons](#), [Information Security Commons](#), [National Security Law Commons](#), [OS and Networks Commons](#), [Other Computer Sciences Commons](#), and the [Social Control, Law, Crime, and Deviance Commons](#)

Scholarly Commons Citation

Blyth, Andrew, "Analysis of Data Erasure Capability on SSHD Drives for Data Recovery" (2018). *Annual ADFSL Conference on Digital Forensics, Security and Law*. 7.
<https://commons.erau.edu/adfsl/2018/presentations/7>

This Peer Reviewed Paper is brought to you for free and open access by the Conferences at Scholarly Commons. It has been accepted for inclusion in Annual ADFSL Conference on Digital Forensics, Security and Law by an authorized administrator of Scholarly Commons. For more information, please contact commons@erau.edu.

EMBRY-RIDDLE
Aeronautical University™
SCHOLARLY COMMONS

(c)ADFSL



ANALYSIS OF DATA ERASURE CAPABILITY ON SSHD DRIVES FOR DATA RECOVERY

Professor Andrew Blyth

Technology Research Centre (TRC), DPG, 50 Brook Street, Mayfair, London, UK.
ab@dpgovernance.com

ABSTRACT

Data Protection and Computer Forensics/Anti-Forensics has now become a critical area of concern for organizations. A key element in this is how data is sanitized at end of life. In this paper, we explore Hybrid Solid State Hybrid Drives (SSHD) and the impact that various Computer Forensics and Data Recovery techniques have when performing data erasure upon a SSHD.

Keywords: Computer Forensics, Data Recovery, Solid State Hybrid Drives (SSHD)

1. INTRODUCTION

The modern storage environment is rapidly evolving. Data may pass through multiple organizations, systems, and storage media in its lifetime. The pervasive nature of data propagation is only increasing as the Internet and data storage systems move towards a distributed cloud-based architecture. As a result, more parties than ever are responsible for effectively sanitizing media and the potential is substantial for sensitive data to be collected and retained on the media. This responsibility is not limited to those organizations that are the originators or final resting places of sensitive data, but also intermediaries who transiently store or process the information along the way. The efficient and effective management of information from inception through disposition is the responsibility of all those who have handled the data. Studies have shown that residual data can still be found on second hand disk drives

(Jones etc al., 2009). Data Protection and GDPR (Linder, 2016) now legally require the report of a data breach, thus the processing of data, and data erasure, at end of life has become of critical concern. The role and function of this paper is to examine the current state of play in relation to the data sanitization of Solid-State Hybrid Drives (SSHD). A number of data recovery methods will be used to explore the capabilities of COTS and Commercial data sanitization methods and they relate to SSHD. The following is the list of data sanitization methods that will be evaluated:

- Standard Open Source Tools.
- Data Sanitization methods such as NIST 800-88 (Kissel etc al., 2014), British HMG Infosec Standard 5, ICT Security Manual (Australian Government, 2016).


The followings is the list of data recovery

methods that will be used to perform the evaluation.

- Non-Invasive/Non-Destructive - Commercial Forensic Tools such as Encase (Widup, 2014), and Access Data FTK (Carbone, 2014), and open-source tools.
- Invasive/Non-Destructive - Commercial Data Recovery Tools such as PC3000 UDMA/SSD and Salvation Data.
- Invasive/Destructive - Chip-Off Data Recovery Tools such as PC3000 Flash and VNR, and attack methods such as JTAG and SPI.

The basic experimental method that will be used is that structured/known-data will be written to every logical block address on the SSHD via a defined and repeatable method. The data sanitization tools will be executed over a sample set and that output shall be analyzed using the data recovery methods listed above. Another key element to be explored in this study relates to the operating system, given the nature of the interface, can the operating system identify a SSHD and then access various elements of it directly? Solid-state hybrid drive (also known by the initialism SSHD) refers to products that incorporate a significant amount of NAND flash memory into a hard disk drive (HDD), resulting in a single, integrated device. The term 'SSHD' is a more precise term than the more general hybrid drive, which has previously been used to describe SSHD devices and non-integrated combinations of solid-state drives (SSDs) and hard disk drives (Michelsoni, 2016). The fundamental design principle behind SSHDs is to identify data elements that are most directly associated with performance (frequently accessed data, boot data, etc.) and store these data elements in the NAND flash memory. This has been shown to be effective in delivering

Figure 1. Basic SSHD Architecture



significantly improved performance over the standard HDD. In Figure-1 we can see the basic SSHD architecture.

The basic principle behind a SSHD is that it functions as a single device with a SATA interface. As far as the host computer system is concerned, the SSHD functions as a single ATA device and conforms to the ATA technical specification (The Serial ATA International Organization, 2016). For the purposes of this study a Seagate Laptop ST500LM000 SSHD has been selected as the test subject. The Seagate Laptop ST500LM000 SSHD is a standard SSHD drive that complies with the SSHD architecture and ATA protocol standards. Making decisions about which data elements are prioritized for NAND flash memory is at the core of SSHD technology. The SSHD using the Solid State component of the drive as a buffer to which output is written to by the controller. The aim is for the Solid State component of the SSHD drive to function as a cache. Solid State Hybrid Drives can operate in one of two modes:

- Self-optimized mode
 - In this mode of operation, the SSHD works independently from the host operating system or host device drives to make all decisions related to identifying data that will be stored in NAND flash

memory. This mode results in a storage product that appears and operates to a host system exactly as a traditional hard drive would.

- Host-optimized mode (or host-hinted mode)
 - In this mode of operation, the SSHD enables an extended set of SATA commands defined in the so-called Hybrid Information feature, introduced in version 3.2 of the Serial ATA International Organization (SATA-IO) standards for the SATA interface. Using these SATA commands, decisions about which data elements are placed in the NAND flash memory come from the host operating system, device drivers, file systems, or a combination of these host-level components (The Serial ATA International Organization, 2016).

2. TECHNICAL ANALYSIS OF SEAGATE LAPTOP SSHD

The Seagate Laptop ST500LM000 SSHD (Model ST500LM000-1EJ162), has 16,383 cylinders, 16 heads and 63 sectors. We can identify the drive's parameters using the *hdparam* tool.

```
$ hdparam -I /dev/sdb
/dev/sdb:
ATA device , non-removable media
Model Number:
ST500LM000-1EJ162
Serial Number:      W3705BXW
Firmware Revision:  DEM3
```

```
Transport:
SATA 1.0a, SATA II Extensions ,
SATA 2.5, SATA 2.6, SATA
3.0
. . . .
Configuration:
Logical      max      current
cylinders    16383   16383
heads        16      16
sectors/track 63      63
—
CHS current  sectors:
16514064
LBA   user  sectors:
268435455
LBA48 user  sectors:
976773168
Logical Sector size:
512 bytes
Physical Sector size:
4096 bytes
Logical Sector -0 offset:
0 bytes
device size :
476940 MBytes
. . . .
```

Analysis of the above shows that the drive is identifying itself as a standard HDD. The next test that we can conduct on the SSHD is to detect the presence of a HPA. The following shows that the HPA is disabled and thus not present on the SSHD.

```
$ hdparam -N /dev/sdb
/dev/sdb:
max sectors: 976773168/976773168
HPA is disabled
```

We can detect the presence of a the DCO via the following:

```
$ hdparam —dco-identify /dev/sdb
/dev/sdb:
```

```

DCO Checksum verified.
DCO Revision: 0x0002
The following features can be
selectively disabled via DCO:
  Transfer modes:
    mdma0 mdma1 mdma2
    udma0 udma1 udma2 udma3 udma4
    udma5 udma6
  Real max sectors: 976773168
  ATA command/feature sets:
  SMART self_test error_log HPA
. . . .

```

The ATA interface supports LBA-48 commands and MAX_LBA is 976.773,168. As far as Technical components are concerned, the Seagate Laptop ST500LM000 SSHD makes use of the following items:

- An LSI B69002V0 drive controller. The motor is controlled by a Texas Instruments SH6966 motor driver.
- A WinBond 25Q80BW16 64MB DDR2 IC acts as the cache for the drive that holds the firmware. This IC is of type 25Q80, with BW packaging.
- To manage flash memory, an eASIC/Seagate 50415 is utilized.
- The 8GB of Toshiba TH58TEG6D2HBA46 BGA/132 NAND flash is present for the "solid state" portion of the SSHD.

The Seagate Adaptive Memory acts as a large cache for frequently accessed data to increase performance, such as booting Windows or loading programs. From the user's perspective, the caching operates in the background, and is designed to work seamlessly and be invisible to the user. Analysis of the PCB showed that from the perspective of data recovery there are two memory integrated circuits worthy of investigation:

- SPI FLASH - WINBOND / 25Q80BW16.
- NAND Flash Solid State Memory - TOSHIBA / TH58TEG6D2HBA46.

3. THE BENCH MARK

To create a single benchmark against which all data sanitization methods could be evaluated, a set of Seagate Laptop ST500LM000 drives had data placed on them. Each drive was connected to a Unix Machine running CentOS 6.9. Using the following a fixed string of **0x5A, 0x5A, 0x5A, 0x5A, 0x5A, 0x5A, 0x5A, 0x5A** was written to every LBA on the device a number of time, thus complying with NIST 800-88 (Kissel et al., 2014). The mount point for the ST500LM000 was identified using the following:

```

[93531.438391] scsi 2:0:0:0
  DirectAccess ATA ST500LM000-1EJ16
  SM15 PQ: 0
[93531.438639] sd 2:0:0:0 [sdb]
  976773168 512-byte logical blocks
[93531.438642] sd 2:0:0:0 [sdb]
  4096-byte physical blocks
[93531.438886] sd 2:0:0:0 [sdb]
  Write Protect is off
[93531.438889] sd 2:0:0:0 [sdb]
  Mode Sense: 00 3a 00 00
[93531.439030] sd 2:0:0:0 [sdb]
  Write/Read cache: enabled.
[93531.439748] sdb: sdb1
[93531.440366] sd 2:0:0:0 [sdb]
  Attached SCSI disk

```

According to the NIST Special Publication 800-88 Section 2.3 (p. 6): "Basically the change in track density and the related changes in the storage medium have created a situation where the acts of clearing and purging the media have converged." That is, for ATA disk drives manufactured after 2001

(over 15 GB), clearing by overwriting the media once is adequate to protect the media from both keyboard and laboratory attack. Please note that this method for placing data on the device does not write data to the DCO and HPA, nor does it write data to the firmware.

```
centos@ajcblyth$ yes "ZZZZZZZZ"
> /dev/sdb 2> /dev/null
```

Data is placed on the device via the following set of LBA 48 Mode write commands.

- The WRITE SECTOR(S) command (OP Code: 0x30h)
 - This command writes from 1 to 256 logical sectors as specified in the Count field. A count of 0 requests 256 logical sectors:
- The WRITE SECTOR(S) EXT command (OP Code: 0x34h)
 - This command is mandatory for devices that implement the 48-bit Address feature set. This command writes from 1 to 65,536 logical sectors as specified in the Count field. A sector count value of 0 requests 65,536 logical sectors.

To validate that data had been written correctly to every Logical Block Address on the device, Encase (Widup, 2014) and FTK (Carbone, 2014) were used to image the SSHD and check that on the string **0x5A, 0x5A, 0x5A, 0x5A, 0x5A, 0x5A, 0x5A, 0x5A** was present.

4. TOOL TECHNICAL ANALYSIS

4.1 Tool: Open Source

To perform data erasure on the SSHD using open source tools the **dd** command was uti-

lized. The following command was used to place zeros on the user space on the SSHD.

```
$ dd if=/dev/zero of=/dev/sdb
```

The process for erasing the data on the SSHD was that the driver was connected to the UNIX system and the above command was executed.

4.1.1 Method:

Non-Invasive/Non-Destructive

The Non-Invasive/Non-Destructive method for analyzing the SSHD is to forensically image the device using tools such as ENCcase (Widup, 2014), and FTK (Carbone, 2014). In addition, a set of open source tools such TAFT (Altheide, 2011) / (Vidstrom, 2015). When both of these techniques were applied to the SSHD they identified nothing but zeros on the User Data area (0 to MAX.LBA) on the drive.

4.1.2 Method:

Invasive/Non-Destructive

This method involves using standard data recovery tools to extract data from the device. These tools two data recovery methods are executed:

- The first method is data extraction via reading data from 0 to MAX.LBA. This method involves the data recovery tools having direct access to the ATA bus and generating/executing their own ATA commands.
- The second method uses standard data recovery and a serial connection to the device to read the firmware located on the device.

Investigation of the SSHD using tools such as PC3000, etc., showed that the NAND Flash memory element of the SSHD is not visible, and thus the SSHD functioned as a single drive, all that was visible was the

Firmware, SMART logs and User Data from 0 to MAX_LBA. Both methods listed above showed that there was no residual data left in either the Firmware/SMART-Logs or the User Data area from 0 to MAX_LBA.

4.1.3 Method: Invasive/Destructive

The Invasive/Destructive method for analyzing the SSHD is to physically remove the **WINBOND / 25Q80BW16** and **TOSHIBA / TH58TEG6D2HBA46** integrated circuits and then using advanced data recovery tools to attempt to read them. All such tests yield negative results, in that no test data that had been written to the drive could be located.

5. CONCLUSIONS

The technical analysis of the SSHD drives architecture (See Figure 1) shows that the SSD component of the SSHD drive is functioning as an inline buffer to the HD component of the SSHD. From a data recovery perspective, the act of writing controller data from Zero to MAX_LBA purges the buffer and overwrites every data element/LBA of the SSHD. While it is true that no data could be recovered from the two integrated circuits (**WINBOND/25Q80BW16** & **TOSHIBA/TH58TEG6D2HBA46**), it should be noted that open source data erasure tools do not remap the drive in-terms of resetting the P/G lists. Thus if an LBA was to fail while holding any test data, the open source tools would not erase this data and thus technically data could still be present on the SSHD.

REFERENCES

- Altheide, C. and Carvey, H. (2011), 'Digital Forensics with Open Source Tools,' *Synpress*.
- Australian Government: Department of Defense: Strategy, Police and Intelligence. (2016), 'Australian Government: Information Security Manual - Controls.'
- Carbone, F. (2014), 'Computer Forensics with FTK,' *Packet Publishing*.
- Jones, A, Dardick, G, Davies, G, Sutherland, I, Valli, C. (2009), 'The 2008 Analysis of Information Remaining on Disks Offered for Sale on the Second Hand Market,' *Journal of International Commercial Law and Technology*, Vol 4, No. 3.
- Kissel, R, Regenscheid, A, Scholl, M, Stine, K. (2014), 'Guidelines for Media Sanitization,' *NIST Special Publication 800-88*.
- Linder A. 2016, 'European Data Protection Law: General Data Protection Regulation 2016,' *CreateSpace Independent Publishing*
- Micheloni, R. (2016), 'Solid State Drives (SSDs) Modeling: Simulation Tools & Strategies,' *Springer*.
- The Serial ATA International Organization. (2016), 'Serial ATA Revision 3.2 Specification.'
- Vidstrom, A. (2015), 'Computer Forensics and the ATA Interface,' *Technical Report, Swedish Defense Research Agency, FOI-R-1638-SE*.
- Widup, S. (2014), 'Computer Forensics and Digital Investigation with EnCase Forensic V7,' *McGraw-Hill Education*.