



Data Recovery from Solid-State Drives

Authors:
Scott Holewinski, Gillware Inc.
Gregory Andrzejewski, Gillware Inc.

12/1/2009

Introduction

The solid-state drive (SSD) industry has an opportunity to address the issue of data loss and recovery from failed SSD devices relatively early in the market and product development cycle. The elimination of moving parts in SSD's should increase the mean time between failure when compared to hard disk drives (HDD's), however still maturing technology and unpredictable operating conditions is already resulting in SSD failure. A certain percentage of these failures will involve the loss of critical data and require data recovery services. The paradigm shift from magnetic to semiconductor-based storage required the development of a completely new set of data recovery techniques. These techniques produce varying degrees of success and are expensive and time consuming to perform. In addition, certain implementations of SSD technologies have been identified that can complicate the recovery process and adversely affect the ability to recover data. By choosing to take a proactive approach and assisting data recovery professionals, the SSD industry will help to ease public concern and increase data recovery success rates while minimizing recovery costs and turn-around times.

Cost Breakdown of Data Recovery

Many factors impact the cost of data recovery from failed storage devices, including equipment, facilities, and human resource expenditures. However, research and development is the biggest contributor to the relatively high price of data recovery.

HDD's and SSD's are incredibly sophisticated devices with multiple potential failure points. Each failure mode requires different techniques in order to recover data stored on the device. The research and development time required to establish reliable and cost-effective recovery procedures for each specific drive and failure mode is substantial. This work is generally performed by experienced teams of electrical and mechanical engineers and computer scientists. Hundreds of new models are released every year, and drive manufacturers are continuously pushing the envelope in terms of performance and capacity. As a result, successful data recovery organizations must invest enormous resources in research and development, with sometimes hundreds of hours spent on the development of a single new technique. Taking the time in the R&D phase to develop efficient data recovery tools and techniques usually results in lower average data recovery costs to the consumer. More specifically, reducing the amount of time spent by an engineer or technician to perform the recovery reduces the cost of the recovery.

Faster turn-around times also mean that the value of the data to the consumer is preserved. In most data recovery scenarios there exists an inverse relationship between the value of the data and the time it takes to recover the data. In other words, the data is never more valuable than the instant it is lost. As potential sales are missed, payrolls come and go, and projected deadlines delayed, the once critical data becomes less important as it is naturally recreated. Therefore, for data recovery to make economic sense the recovery process must be accomplished both quickly and cost-effectively. Most data recovery professionals agree that except in cases where data cannot be

recreated, there is a precipitous drop-off in the number of customers willing to pay for their lost data when recovery times exceed three weeks. Figure 1 depicts the delicate balance that exists in the data recovery industry between the value of the lost data to the consumer, and the cost and turn-around time to perform the recovery.

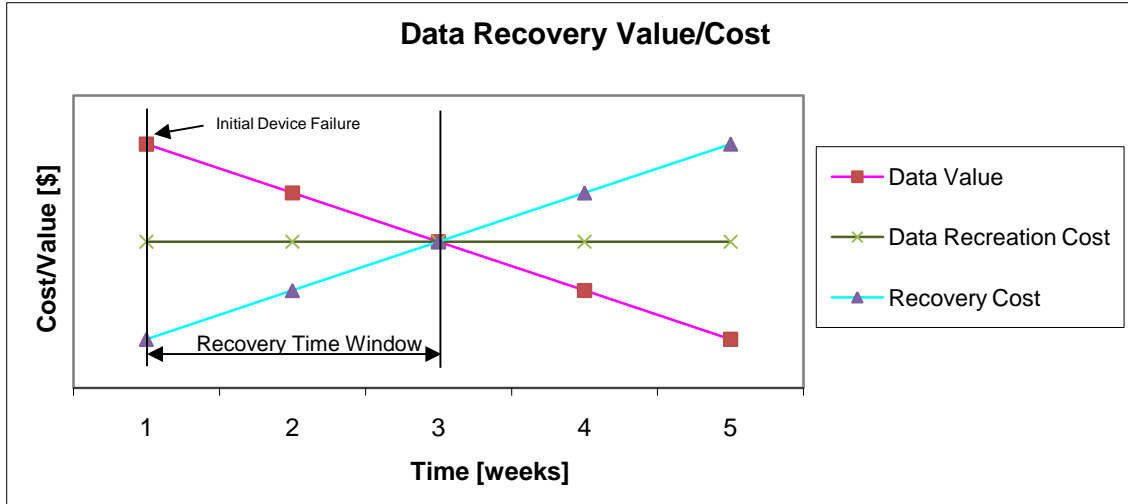


Figure 1. A graphical representation of the relationship between the value of lost data to the consumer and the cost of the recovery.

Through a commitment to R&D Gillware Inc. has been able to significantly reduce the turn-around time and total cost for a single HDD data recovery from an industry average of \$1500 and three weeks respectively. For the fiscal year 2009 the average HDD data recovery at Gillware Inc. cost \$694 and took six business days to complete, staying well within the recovery time window shown in Figure 1.

Years of experience and well-defined techniques have stabilized the average cost and turn-around time for data recovery from HDD's. SSD data recovery on the other hand is a discipline that is being developed as the SSD technology grows. As a result, the cost and recovery time from SSD's can vary dramatically depending on the manufacturer and specific parameters of the device. Solid-state storage technology represents an entirely new set of engineering problems to research teams at data recovery organizations. SSD manufacturers are pushing the technology envelope in order to increase drive storage capacities while attempting to improve device reliability. The result is a blistering pace of change with frequent releases of new designs. Each new design represents new firmware, different wear leveling algorithms and controllers, and revised PCB layout. Staying ahead of the SSD data recovery curve is a challenge, keeping in mind that the delicate balance between recovery turn-around times and cost must be maintained. Although SSD data recovery techniques are progressing, they lag behind the streamlined and efficient procedures used to recover data from HDD's. As a result, the average SSD recovery at Gillware Inc. costs \$2850 and takes approximately three weeks to perform. The data recovery cost discrepancy between SSD's and HDD's is shown in Figure 2.

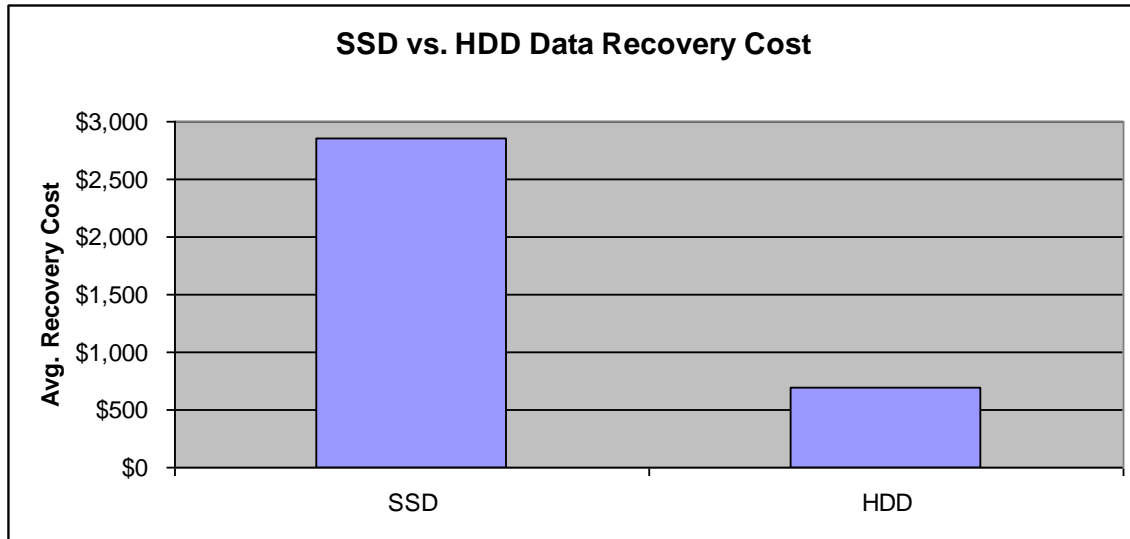


Figure 2. Cost of data recovery performed on SSD's vs. HDD's.

For data recovery from SSD's to be a viable option for the growing SSD market, the cost and turn-around time must be brought in-line with those of HDD's.

Device Failure and the Customer Relationship

There is a fine line that every electronic storage device manufacturer must walk when dealing with device failure. The fact that devices will fail is a virtual certainty.

Regardless of the reason for failure, manufacturer defect or abuse by the end user, the consumer expects the device manufacturer to provide a certain level of assistance in recovering the stored electronic data. For some queues on how these situations can be handled one need look no further than the approaches implemented by hard disk drive OEM's and computer manufacturers.

Each hard disk drive and computer manufacturer has a different approach to handling data recovery situations resulting from HDD failure. These approaches range from an apology and an offer to replace the device if it's under warranty, to providing in-house data recovery services paid for by the customer. Although both are viable options, neither is popular with consumers, proving to be public relationship missteps when utilized. The most popular approach is to provide the customer with a short list of data recovery providers that have been vetted by the OEM's as capable and professional data recovery organizations. In exchange for being placed on the list the data recovery providers commonly offer the customer a small discount, and the HDD manufacturers give the data recovery providers a small amount of technical assistance when necessary. By ensuring that the customer's at the very least have a positive data recovery experience, OEM's are able to lessen the potentially negative impact device failure can have on future sales.

HDD Data Recovery Process

The HDD data recovery process can be broken down into four phases: drive failure analysis, drive restoration, drive imaging, and data extraction. Although all four phases represent unique challenges in order to maintain process efficiency, drive failure diagnosis and restoration are where the majority of the engineering resources are required. HDD's have three primary failure modes: logical, electrical, and mechanical. Mechanical failures are largely isolated to the read/write head assembly or the spindle motor, and are usually the result of mechanical fatigue or environmental abuse i.e. HDD is knocked over while running. Electrical failures can be caused by numerous conditions, but most are the result of power surges or individual circuit component failure. Logical failure usually consists of corruptions to the firmware area of the HDD and can happen naturally or as the result of an upstream failure such as an intermittent read/write head.

Complicating the drive failure diagnosis process is the interdependency of the three failure categories. For example, it is not uncommon for a shorted control board to cause a read/write head failure, or for an intermittent read/write head to corrupt the firmware zone. As a result, a certain level of experience and intuition is required by the engineers tasked with diagnosing the failure mode of the device. These engineers are highly compensated individuals and account for a significant portion of the total cost of the recovery. After identifying the root cause of the drive failure recovery technicians can perform the necessary repair work required to restore the drive back to a functional state. A proper diagnosis of the failure mode results in lower engineering costs and faster turn-around times.

The case study in the following section outlines the steps that were followed in an actual HDD data recovery of a failed 1 TB hard drive.

HDD Data Recovery Case Study

Failure Description: Drive is no longer spinning up. No unusual sounds heard from the HDD.

Initial Failure Diagnosis: Starting with an investigation of the control board a distinct "burn" smell was noticed coming from the PCB. Further inspection identified four failed control board components. Figure 3 is a picture of the particular area of the control board that was electrically damaged.




Figure 3. Shown is a shorted diode that has “popped” after a power surge.

Following inspection of the HDD control board, an internal test of the read/write head assembly was performed. A non-invasive electrical test of the HDD head assembly identified that at least one of the eight read/write heads was no longer functioning. At this point the drive was taken into Gillware’s ISO 5 certified Class-100 cleanroom for further analysis. Figure 4 shows the physical condition of the read/write heads.

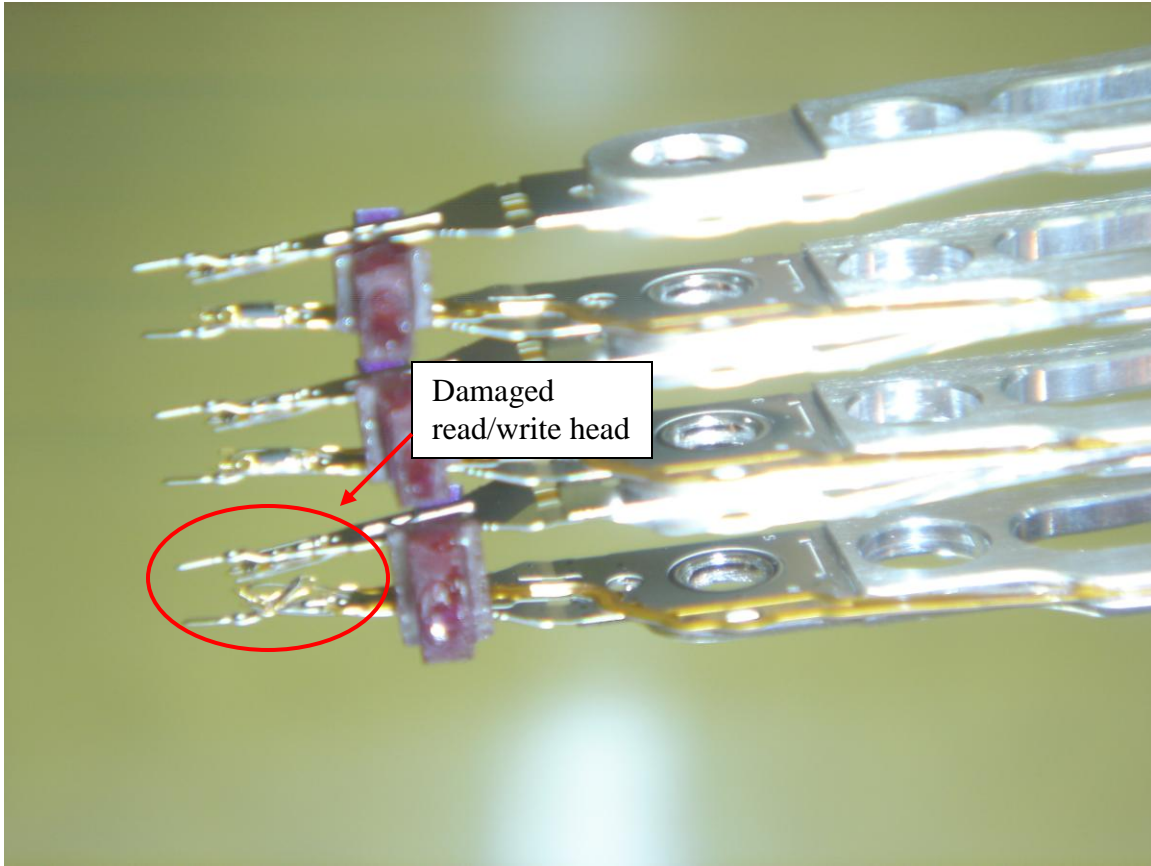


Figure 4. Read/write head damage resulting from sudden power loss.

Final Failure Analysis: Drive has two failures preventing it from properly functioning. A power surge has taken out one or more components on the HDD control board. The sudden loss of power prevented the heads from parking properly, instead they were stuck platter surface. This resulted in the bent read/write head shown in figure 4.

Drive Restoration: Restoration of the hard drive begins with repair work on the failed electrical components on the control board. Following the control board rework, the damaged read/write heads are carefully replaced. Figure 5 shows the open HDD chassis with new read/write heads installed.




Figure 5. Open HDD chassis with healthy read/write heads installed.

Read/write head compatibility is a major issue on most modern hard drives. It is quite rare to find a set of replacement heads that immediately work when transplanted from one HDD and placed in another. Therefore, following head stack replacement Gillware engineers use proprietary logical tools to restore the drive to an operational state. Prior to HDD imaging, areas of minor platter damage are addressed in order to prevent further damage during the imaging process.

Drive Imaging: Following drive restoration the HDD is moved to the drive imaging phase of the recovery process. Drive imaging involves making a direct byte-for-byte copy of the HDD for use in the logical processing and data extraction phase of the recovery. The amount of time required to image a drive can vary drastically depending on the specific details of the recovery case. The 1 TB HDD in this case study with a repaired control board and replaced read/write heads took approximately 27 hours to image, completing with a 99.99% read of the available sectors on the HDD. The reduced drive performance is the result of adaptives deviation; a symptom caused by the non-native control board and read/write heads. It is not uncommon for HDD's to take multiple days to image if the damage to the magnetic media is severe.

Data Extraction: The image copy is provided to the logical engineers where they are tasked with recreating the file system and extracting user-data. The HDD in this case

study had only electrical and mechanical damage. This fact coupled with a very good image copy of the drive means that there was no logical corruption to the file structure. When this is the case the data extraction process at Gillware is automated through the use of proprietary extraction tools. After the data is extracted, the customer verifies the recovered data via the Gillware File Viewer application and the data is transferred to a new external transfer drive for delivery to the customer.

Summary: The recovery of the 1 TB HDD in this case study was a grade-A recovery. All user data was recovered intact and fully functional. The total in-lab time for the recovery was approximately 6 business days. The total engineering time was 5 hours [0.75 hours for evaluation, 3 hours for drive restoration, 1.25 hours for logical processing and extraction]. The total cost for the recovery was \$1000 [\$875 for the recovery, \$125 for a new 1TB external transfer drive to ship the data back to the customer].

SSD Data Recovery Process

Solid-state devices share many of the same failure modes exhibited by HDD's. Since SSD's are a direct replacement for HDD's in most applications and are subject to many of the same stresses, some SSD failure modes are similar to those of HDD's. The most significant difference between the two technologies is an SSD's lack of moving. As a result, SSD's have no instances of mechanical failure. Shared failure modes aside, the techniques and processes for recovering data from the two storage technologies differ greatly.

SSD's afford data recovery professionals opportunities to recover data from failed devices not available with HDD's. The Holy Grail of HDD data recovery is a device that can read HDD platters independent of the hard drive. Although accomplished in laboratory environments with varying degrees of success, this technique is not a viable option for the recovery of large amounts of data commonly stored on modern drives. The process is simply too slow, and requires too much user input to make economic sense. Drive restoration is a more efficient and cost-effective approach. SSD's on the other hand store data in non-volatile memory chips that can be easily read independent of the device that originally wrote the data. This presents the possibility of an alternate recovery process for SSD's where no repairs are necessary. Start by individually imaging each memory chip, assemble the individual chip images into a single drive image, and extract the data.

The reconstruction of a single drive image is the most time consuming and costly aspect of the SSD data recovery process. With HDD data recovery the end result of the drive imaging process is a single complete image starting with sector zero and ending with the last sector on the HDD. Compare this to SSD's where the output of the imaging phase of the recovery process is N individual chip images resulting from reading N number of chips on the device, i.e. an SSD with 16 memory chips produces 16 individual chip images. These images must be reconstructed into a single device image prior to

proceeding with the data extraction process. The drive reconstruction phase is the most demanding stage of the process, as the method used to spread data across each memory chip varies from model to model. With no information about how the data is striped across the memory chips comprising the full array, the only options is to manually find key file structure indicators, and use those indicators to reassemble the data.

Some of the drawbacks to the independent memory chip imaging data recovery approach are evident in the following case studies. All three have slightly different hardware, software, and end-user implementations. The impact these different implementations have on the data recovery process is illustrated.

SSD Data Recovery Case Study 1:

Summary Result: Successful recovery

SSD Details: 128 GB SSD with 16 (8 GB) TSOP48 memory chips. No after-market or factory direct encryption.

Failure Description: Computer does not recognize the SSD

Initial Failure Diagnosis: The SSD appears pristine; there is no evidence of electrical damage, yet the device shows no response when connected to a host.

Final Failure Analysis: The drive is suffering from a logical failure, likely due to firmware corruption. At this time, no tools for firmware repair are available and the only option for a timely successful recovery is to read the contents of each of memory chip and reconstruct the drive image.

Memory Chip Read: Each chip is removed from the SSD and its contents copied to a file on a PC. Figure 6 shows the reading of one of the 16 chips on this SSD device.




Figure 6: Reading of an 8 GB memory chip.

Drive Image Reconstruction: Without knowledge of how the SSD keeps track of the storage spread across each memory chip, file system structures are used to reconstruct the disk image. Regardless of the file system used, there will usually be partition table found at the first logically-addressable sector. The customer reported the system was running Windows XP, so a Master Boot Record (MBR) should be found sector 0. Each of the 16 reads is searched for the signature of an MBR. After the MBR is located, the mapping of data physical location to logical sector can proceed.

As suspected, there are two partitions on this SSD: a small, FAT16 system-restore partition and a large NTFS partition. The MBR provides the location and size of each partition and the reads are searched for the corresponding boot sectors. Each boot sector must be located at the logical start of the partition, indicated by the MBR.

The FAT16 file system has a lengthy list of mostly sequential values called a FAT Table immediately following the boot sector. The FAT Table provides a lot of useful information because, due its relatively sequential nature, the first value in the following logical sector can be predicted and the reads are searched to build this table. Using the physical location of successive logical sectors, a pattern begins to appear as to how the data is organized throughout the reads that can be used to build the rest of the disk image.

Data Extraction: Once the disk image is built, data extraction can proceed in the same manner as with an HDD.

Detailed Final Result: The result of the 128 GB SSD in this case study was a grade-A recovery. All user data was recovered in tact and fully functional. The total in-lab time for the recovery was approximately 2.5 weeks. The total engineering/machine time was 22 hours [2 hours for evaluation, 8 hours for de-soldering and reading memory chips, 12

hours for logical processing and data extraction]. The total cost for the recovery was \$3000.

SSD Data Recovery Case Study 2:

Summary Result: Unsuccessful recovery

SSD Details: 128 GB SSD with 16 (8 GB) TSOP48 memory chips. Aftermarket full-disk encryption implemented.

Detailed Final Result: The standard SSD data recovery procedure outlined in case study 1 was followed. Images of all 16 individual memory chips were generated by desoldering the chips and reading them on a TSOP48 fixture. Following chip imaging the recovery specialist began the image reconstruction process. As in case study 1 the technician attempted to identify common file structure indicators in order to identify the striping of the data across the 16 images. However, the drive in this case study was from a large enterprise customer that implements full-disk encryption on all company computers. As a result, no file structure indicators could be identified and the data reconstruction procedure failed. No data could be recovered in a timeframe acceptable to the customer.

Additional Comments: Full-disk encryption is becoming increasingly more common among Gillware's enterprise client base. Customers often ask Gillware's assistance when sourcing encryption products in order to avoid data recovery complications in the future. Most encryption vendors provide tools for decrypting a drive image with proper credentials, which Gillware uses in the recovery process. The merits of these tools often influence Gillware's advice to clients. Unfortunately these tools only succeed when given a complete, correct image copy. With no knowledge of how to reassemble the individual memory chip images back into a single disk image, Gillware technicians are unable to recover data.

Potential Future Solution: Gillware hopes to partner with SSD manufacturers in order to help enterprise customers recover data in situations where full-disk encryption is utilized. Formal partnerships will allow for the protected sharing of sensitive and proprietary technical information about each SSD. With detailed knowledge of the device's Flash Translation Layer, firmware, controller, and ECC implementation, Gillware technicians will no longer need to rely on file system structures and will be able to successfully recover data from SSD's with full-disk encryption. The end result of these partnerships will be higher SSD recovery success rates and the preservation of relationships with key enterprise customers for both SSD manufacturers and Gillware alike.

SSD Data Recovery Case Study 3:

Summary Result: Unsuccessful recovery

SSD Details: 128 GB SSD with 16 (8 GB) TSOP48 memory chips. Full-disk hardware-level encryption of data stored on memory chips.

Detailed Final Result: Similar to case study 2, no file structure indicators were discernible as a result of the data being encrypted by SSD device. With no knowledge of the Flash Translation Layer (FTL) or the manner in which the encryption is performed, Gillware technicians are unable to recover any data.

Potential Future Solution: Offering storage devices with hardware-level encryption can be a powerful marketing tool. This is especially true when looking to land lucrative enterprise contracts where encryption is an absolute requirement. There are two issues that Gillware encounters when dealing with storage devices implementing full-disk encryption. First, many end-users are unaware of the encryption and are frustrated when they discover the negative impact encryption technology can have on the recovery process. Second, without vendor support, Gillware has no means of decrypting the data.

Successful recovery of data from SSD's with hardware-level encryption requires both an understanding of the FTL and means to decrypt the drive image. Depending on how encryption keys are generated and stored, Gillware suggests maintaining a protected database of keys for use in extreme case where the key cannot be directly obtained from the failed device. As an alternative, SSD manufacturers may choose to provide emulation and decryption software tools to data recovery partners, similar to those provided by software-based encryption vendors. Critical points for future discussions include the formation of partnerships between SSD manufacturers and data recovery providers that both protect intellectual property while still allowing data recovery to be performed successfully.

The Future of SSD Data Recovery

Gillware currently has two primary organizational objectives focused on dealing with the issues surrounding SSD data recovery. The first is the development of reliable data recovery tools and techniques that allow engineers to recover data from failed SSD devices. The second is to eliminate the large discrepancy in cost and turn-around time that exists between HDD and SSD data recovery. Both objectives are closely linked and must be accomplished in unison. That is, the ability to recover data from failed devices is inconsequential if the recovery cannot be performed in a cost-effective and efficient manner.

The R&D team at Gillware Inc. is working hard to meet both objectives and become the industry leader in SSD data recovery. However, as the SSD case studies point out,

improvements are required in order to decrease the amount of engineering time necessary to perform the recoveries and to improve recovery success rates. A reduction in engineering time will directly correlate to lower overall recovery costs and improved success rates help to bolster customer satisfaction.

Solid-state storage devices offer many advantages over HDD's when looking at data recovery techniques and procedures. The most significant of these advantages is the ability to read data from individual memory chips independent of the host device. This recovery technique will eventually lead to better overall success rates and potentially lower data recovery costs when compared to data recovery from HDD's. However, certain obstacles must first be overcome before these benefits are made a reality.

The reading of the individual memory chips is an essential and inefficient step in the SSD data recovery process. Most SSD memory chips come in a TSOP48 package which requires the use of a specialized fixture (Figure 6) to be read with a commercial device programmer. These fixtures are quite effective when used with new, pristine devices. However, ensuring proper electrical contact with the pins of a device unsoldered from a PCB is a constant struggle. Pins can easily be bent if the IC is not removed carefully. Furthermore, any residual bits of solder can affect the delicate alignment of the pin-contact fingers. It remains to be seen whether this situation will improve with the transition to ball grid array packages. Gillware engineers are currently working on a solution that will streamline the memory chip imaging process significantly reducing turn-around times. It is also possible that SSD manufacturers could decrease the need for removing and individually reading the memory chips by implementing certain technology common on HDDs.

Every HDD has a vendor-specific mechanism for manipulating device firmware over the ATA interface and sometimes through other means, such as an undocumented RS-232 or JTAG connection. Gillware engineers are speculating that SSD's might have a similar implementation. If this is the case, SSD's with certain firmware corruptions could potentially be repaired and the device restored to a functional state. If the firmware could not be repaired, the ability to obtain the raw contents of the memory chip would be incredibly valuable. In either situation the need to desolder and individually read each memory chip is eliminated.

Whether dumping the drive image directly from the SSD, or reconstructing it manually, an understanding about how the controller maintains the Flash Translation Layer (FTL) is critical. This logical-sector to physical-address mapping is the heart of any wear leveling implementation and is currently the biggest hurdle Gillware faces with SSD data recovery. In SSD case study 1, Gillware engineers were able to discern enough about the FTL by looking at the physical location of critical file system structures known to reside at a given logical sector. While yielding a successful recovery, this method will not scale to the volume of business Gillware currently does with HDD's. This technique is also unsuccessful in situation involving full-disk encryption. Without assistance from

SSD manufacturers, data recovery providers will struggle to match the success rate levels established with HDD recoveries.

There are six ways by which SSD manufacturers can assist data recovery partners, helping to improve SSD data recovery success rates and reducing costs:

1. Provide technical details of the FTL
2. Supply documentation of vendor-specific ATA commands for firmware and FTL manipulation
3. Allow access to the appropriate cipher in the presence of hardware-level encryption
4. Provide information about the ECC implementation
5. Grant access to data sheets for SSD controllers and non-volatile memory
6. Supply controller emulation tools
7. Co-develop with Gillware engineers improved systems for obtaining memory chip reads

Conclusion

The cost and turn-around times associated with SSD data recovery will improve over the years to come. As SSD technology, standards, and designs stabilize, so will the tools and techniques required to recover data from failed devices. How quickly this happens will depend largely on the level of cooperation that is provided by the SSD industry to key data recovery partners. For these partnerships to be successful, both parties will need to work together to guarantee that sensitive proprietary information is protected. Current SSD data recovery techniques are reactive, primarily being developed as device failures occur and arrive in the lab. This approach to data recovery is affective, but expensive and time consuming to perform. Through a collaborative effort between the solid state and data recovery industries, it is possible to predict and plan for a majority of SSD failures. The benefits to this proactive approach will be lower recovery costs, improved turn-around times, and better overall success rates. Working closely with data recovery professionals has the added benefit of in-the-trenches failure analysis that can be relayed back to device manufacturers. This information can be used by reliability and design engineering groups to improve device reliability, preventing future failures and moving solid-state technology forward.



Gillware Inc.

10 Terrace Ct Suite 103 · Madison, WI 53718 · Phone: 877.624.7206 · www.gillware.com

©Copyright 2009 Gillware Inc. All Rights Reserved.