Recovering Evidence from SSD Drives: Understanding TRIM, Garbage Collection, and Exclusions

Yuri Gubanov and Oleg Afonin

In 2012, *DFI News* published an article called “Why SSD Drives Destroy Court Evidence, and What Can Be Done About It” (www.dfinews.com/articles/2012/09/why-ssd-drives-destroy-court-evidence-and-what-can-be-done-about-it-part-1 [1]). Back then SSD self-corrosion, TRIM, and garbage collection were little known and poorly understood phenomena. In 2014, the situation looks different. Having handled numerous cases involving the use of SSD drives and gathered a lot of statistical data, we now know things about SSD drives that allow forensic specialists to obtain information from them despite the obstacles.

SSD Self-Corrosion

The effect of SSD self-corrosion, as well as the root cause, is well covered by existing publications, including our own 2012 paper on SSD forensics. The evidence self-destruction process is triggered by the TRIM command issued by the operating system to the SSD controller at the time the user either deletes a file, formats a disk, or deletes a partition. The data destruction process is only triggered by the TRIM command; the data destruction itself is carried out by the separtate process of background garbage collection.

In many cases the TRIM command is not issued at all. This article discusses these exclusions to gain a better understanding of the situations when deleted data can still be recovered from an SSD drive.

Deterministic Read After Trim

Experiences recovering information from SSD drives vary greatly among SSD users.

“I ran a test on my SSD drive, deleting 1,000 files and running a data recovery tool five minutes later. The tool discovered several hundred files, but an attempt to recover them returned a bunch of empty files filled with zeroes,” said one Belkasoft customer.

“We analyzed an SSD drive obtained from a suspect’s laptop and were able to recover 80% of deleted files several hours after they’ve been deleted,” said another user.

Why such inconsistency in user experiences? The answer lies in the way the different SSD drives handle trimmed data pages.

Some SSD drives implement what is called Deterministic Read After Trim (DRAT) and Deterministic Zeroes After Trim (DZAT), returning all-zeroes immediately after the TRIM command releases a certain data block, while others do not implement this protocol and will return the original data until it’s physically erased with the garbage collection algorithm.

With non-deterministic TRIM, each read command after a Trim may return different data, while with both DRAT and DZAT, all read commands after a TRIM return the same data.

As we can see, in some cases the SSD will return non-original data (all zeroes, all ones, or some other non-original data) not because the physical blocks have been cleaned immediately following the TRIM command, but because the SSD controller says that there is no valid data held at the trimmed address on a logical level previously associated with the trimmed physical block.

Acquiring Evidence from SSD Drives

So far the only practical way of obtaining evidence from an SSD drive remains traditional imaging (with dedicated hardware/software combination), followed by an analysis with an evidence discovery tool.

Whether or not you’ll be able to access deleted information stored on an SSD drive depends on the system using the SSD drive supporting the full TRIM/garbage collection chain.

Operating System Support

TRIM is a property of the operating system as much as it is a property of an SSD device. Older file systems do not support TRIM. Wikipedia http://en.wikipedia.org/wiki/Trim_(computing) [2] has a comprehensive table detailing the operating system support for the TRIM command. Put simply, TRIM support was only added in Windows version 7 and later. In addition, Windows only supports TRIM if the drive is formatted with the NTFS file system. If the SSD drive carries a different file system such as FAT32 or exFAT, there will be no TRIM on that drive.

In Mac OS X, TRIM is only supported for SSD drives approved and supplied by Apple, and only starting with OS version 10.6.8. Older builds of Mac OS X do not support TRIM. Notably, user-installed SSD drives not supplied by Apple itself are excluded from TRIM support.

Old or Basic SSD Hardware

Not all SSD drives support TRIM and/or background garbage collection. Older SSD drives as well as SSD-like flash media used in basic tablets and sub-notebooks (such as certain models of ASUS Eee) do not support the TRIM command. For example, Intel started manufacturing TRIM-enabled SSD drives with drive lithography of 34nm (G2); their 50nm SSDs do not have TRIM support.

In reality, few SSD drives without TRIM survived that long. Many entry-level sub-notebooks use flash-based storage often mislabeled as “SSD” that does not feature garbage collection or support the TRIM protocol.

External Drives, USB Enclosures, and NAS

The TRIM command is fully supported over the SATA interface, including the eSATA extension, as well as SCSI via the UNMAP command. If an SSD drive is used in a USB enclosure or installed in certain types of NAS storage, the TRIM command will not be communicated via the unsupported interface.

PCI-Express and PCIe SSDs

Interestingly, the TRIM command is not natively supported by any version of Windows for many high-performance SSD drives occupying the PCI Express slot. Do not confuse PCI Express SSDs with SATA drives carrying M.2 or mSATA interfaces. PCI Express boards implementing an SATA port are also excluded from this exclusion.

RAID

The TRIM command is not supported on most RAID configurations (with few rare exceptions). SSD drives working as part of a RAID array can be analyzed.

A notable exception to this rule would be the modern RAID 0 setup using a compatible chipset (such as Intel H67, Z77, Z87, H87, Z68) accompanied by the correct drivers (the latest Rapid Storage Technology driver from Intel allegedly works) and a recent version of BIOS. In these configurations, TRIM can be enabled.

Corrupted Data

Surprisingly, SSD drives with corrupted system areas (damaged partition tables, skewed file systems, etc.) are easier to recover than healthy ones. The TRIM command is not issued over corrupted areas because files are not properly deleted. They simply become invisible or inaccessible to the operating system. Many commercially available data recovery tools (e.g., Intel Solid-State Drive Toolbox with Intel SSD Optimizer, OCZ SSD Toolbox) can reliably extract information from logically corrupted SSD drives.

Bugs in SSD Firmware

Firmware used in SSD drives may contain bugs, often affecting the TRIM

functionality and/or messing up garbage collection. For example, OCZ Agility 3 120 GB shipped with buggy firmware v. 2.09, in which TRIM did not work. Firmware v. 2.15 fixed the TRIM problem, while v. 2.22 introduced issues with data loss on wake-up after sleep. Firmware v. 2.25 fixed that but disrupted TRIM operation again (http://www.overclock.net/t/1330730/ocz-firmware-2-25-trim-doesnt-work-bug-regression-bad-ocz-experience [3]). A particular SSD drive may or may not be recoverable depending on which bugs were present in its firmware.

Bugs in SSD Over-Provisioning

SSD over-provisioning is one of the many wear-leveling mechanisms intended to increase SSD life span. Some areas on the disk are reserved on the controller level, meaning that a 120 GB SSD drive carries more than 120 GB of physical memory. These extra data blocks are called the over-provisioning (OP) area and can be used by SSD controllers when a fresh block is required for a write operation. A dirty block will then enter the OP pool and will be erased by the garbage collection mechanism during the drive's idle time.

In regard to SSD over-provisioning, firmware bugs can affect TRIM behavior in other ways. For example, revealing trimmed data after a reboot/power off. Solid-state drives remap constantly after TRIM to reallocate addresses in the OP pool. As a result, the SSD reports a trimmed data block as writeable (already erased) immediately after TRIM. Obviously, the drive did not have the time to actually clean old data from that block. Instead, it simply maps a physical block from the OP pool to the address referred to by the trimmed logical block.

What happens to the data stored in the old block? For a while, it contains the original data (in many cases it's compressed data, depending on the SSD controller). However, as that data block is mapped out of the addressable logical space, the original data is no longer accessible or addressable.

Sounds complex? You bet. That's why even seasoned SSD manufacturers may not get it right the first time (as discussed on OCZ Technology Forum at http://www.ocztechnologyforum.com/forum/showthread.php?96382-Deterministic-Read-After-Trim [4]). This creates issues when, after deleting data and rebooting the PC, some users would see the old data back as if it was never deleted. Apparently, because of the mapping issue, the new pointers would not work as they should due to a bug in the drive's firmware. OCZ released a firmware fix to correct this behavior, but similar (or other) bugs may still affect other drives.

SSD Shadiness: Manufacturers Bait-and-Switch

When choosing an SSD drive, customers tend to read online reviews. Normally, when a new drive gets released, it is reviewed by various sources soon after it becomes available. The reviews get published, and customers often base their choice on them.

But what if a manufacturer silently changes the drive's specs without changing the model number? In this case, an SSD drive that used to have great reviews suddenly becomes much less attractive. This is exactly what happened with some manufacturers. According to ExtremeTech

(http://www.extremetech.com/extreme/184253-ssd-shadiness-kingston-and-pny-caught-bait-and-switching-cheaper-components-after-good-reviews [5]), two well-known SSD manufacturers, Kingston and PNY, were caught pulling a bait-and-switch with cheaper components after getting good reviews. In this case, the two manufacturers were launching their SSDs with one hardware specification, and then quietly changed the hardware configuration after reviews went out.

So what does this mean for us? Well, the forensic-friendly SandForce controller was found in the second revision of PNY Optima drives. Instead of the original Silicon Motion controller, the new batch of PNY Optima drives had a different, SandForce-based controller known for its less-than-perfect implementation of garbage collection, which left data on disks for a long time after it was deleted.

Small Files: Slack Space

Remnants of deleted evidence can be acquired from so-called slack space as well as from MFT attributes.

In the word of SSD, the term “slack space” receives a new meaning. Rather than being a matter of file and cluster size alignment, “slack space” in SSD drives deals with the different sizes of minimum writeable and minimum erasable blocks on a physical level.

In SSD terms, “page” is the smallest unit of storage that can be written to. The typical page size of today’s SSD is 4 KB or 8 KB.

“Block,” on the other hand, is the smallest unit of storage that can be erased. Depending on the design of a particular SSD drive, a single block may contain 128 to 256 pages.

As a result, if a file is deleted and its size is less than the size of a single SSD data block, or if a particular SSD data block contains pages that still remain allocated, that particular block is not erased by the garbage collection algorithm. In practical terms, this means that files or file fragments (chunks) smaller than 512 KB or 2 MB depending on SSD model, may not be affected by the TRIM command, and may still be forensically recoverable.

However, the implementation of the Deterministic Read After Trim (DRAT) protocol by many recent SSD drives makes trimmed pages inaccessible via standard SATA commands. If a particular SSD drive implements DRAT or DZAT (Deterministic Read Zero After Trim), the actual data may physically reside on the drive for a long time, yet it will be unavailable to forensic specialists via standard acquisition techniques. Sending the SSD drive to the manufacturer might be the only way of obtaining this information on a physical level.

Small Files: MFT Attributes

Most hard drives used in Windows systems use NTFS as their file system. NTFS stores information about the files and directories in the Master File Table (MFT). MFT contains information about all files and directories listed in the file system. In other

words, each file or directory has at least one record in MFT.

In terms of computer forensics, one particular feature of MFT is of great interest. Unique to NTFS is the ability to store small files directly in the file system. The entire content of a small file can be stored as an attribute inside an MFT record, greatly improving reading performance and decreasing wasted disk space ("slack" space).

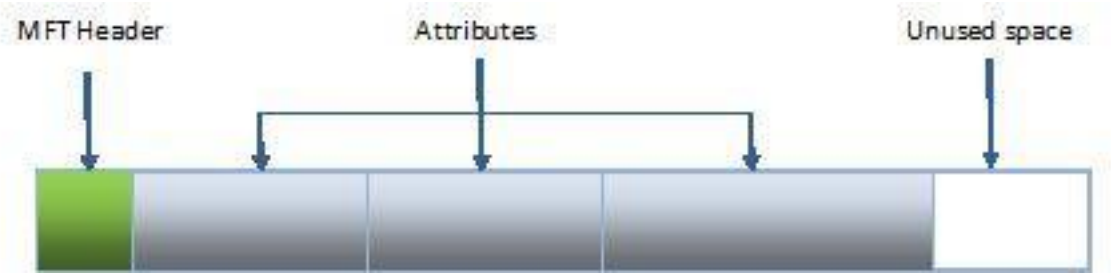

As a result, small files being deleted are not going anywhere. Their entire content continues residing in the file system. The MFT records are not emptied and are not affected by the TRIM command. This in turn allows investigators to recover such resident files by carving the file system.

How small does a file have to be to fit inside an MFT record? Very small. The maximum size of a resident file cannot exceed 982 bytes. Obviously, this severely limits the value of resident files for the purpose of digital forensics.

Reality Steps In: Why Real SSDs are Often Recoverable
In reality, things may look different from what was described above. In our lab we've seen hundreds of SSD drives acquired from a variety of computers. Surprisingly, we were able to successfully carve deleted data from the majority of SSD drives taken from inexpensive laptops and sub-notebooks such as ASUS Eee or ASUS Zenbook. Why? There are several reasons, mainly "cost savings" and "miniaturization," but sometimes it's simply over-engineering.

1. Inexpensive laptops often use flash-based storage, calling that an SSD in their marketing. In fact, in most cases it's just slow, inexpensive, and fairly small flash-based storage having nothing to do with real SSD drives.
2. Ultrabooks and sub-notes have no space to fit a full-size SSD drive. They used to use SSD drives in PCIe form factor (as opposed to M.2 or mSATA) which did not support the SATA protocol. Even if these drives are compatible with the TRIM protocol, Windows does not support TRIM on non-ATA devices. As a result, TRIM is not enabled on these drives.
3. SSD drives are complex devices requiring complex firmware to operate. Many SSD drives were released with buggy firmware effectively disabling the effects of TRIM and garbage collection. If the user has not upgraded their SSD firmware to a working version, the original data may reside on an SSD drive for a long time.
4. The fairly small (and inexpensive) SSD drives used in many entry-level notebooks lack support for DRAT/DZAT. As a result, deleted (and trimmed)

data remain accessible for a long time, and can be successfully carved from a promptly captured disk image.

5. On the other end of the spectrum are the very high-end, over-engineered devices. For example, Acer advertises its Aspire S7-392 as having a RAID 0 SSD. According to Acer marketing, “RAID 0 solid state drives are up to 2X faster than conventional SSDs. Access your files and transfer photos and movies quicker than ever!” This looks like over-engineering. As TRIM is not enabled on RAID SSDs in any version of Windows, this ultra-fast non-conventional storage system may slow down drastically over time (which is exactly why TRIM was invented in the first place). For us, this means that any data deleted from these storage systems could remain there for at least as long as it would have remained on a traditional magnetic disk. Of course, the use of the right chipset (such as Intel H67, Z77, Z87, H87, Z68) accompanied with the correct drivers (the latest RST driver from Intel allegedly works) can in turn enable TRIM. However, we have yet to see how this works in reality. (http://www.anandtech.com/show/6477/trim-raid0-ssd-arrays-work-with-intel-6series-motherboards-too [6])

Conclusion

SSD forensics remains different. SSDs self-destroy court evidence, making it difficult to extract deleted files and recover destroyed information. Numerous exceptions still exist, allowing forensic specialists to access destroyed evidence on SSD drives used in certain configurations.

More SSD drives appear to follow the Deterministic Read After Trim (DRAT) approach defined in the SATA standard set a long time ago. This in turn means that a quick format is likely to instantly render deleted evidence inaccessible to standard read operations, even if the drive is acquired with a forensic write-blocking imaging hardware immediately after deletion.

SSD drives are getting more complex, using over-provisioning support for better performance and wear leveling. However, because of the increased complexity, even seasoned manufacturers released SSD drives with buggy firmware, causing improper operation of TRIM and garbage collection functionality. Considering just how complex today’s SSD drives have become, it’s surprising these things do work, even occasionally.

The playing field is constantly changing, but what we know now about SSD forensics gives hope.

Yuri Gubanov *is a computer forensics expert and a frequent speaker at industry conferences. Yuri is the Founder and CEO of Belkasoft, the manufacturer of computer forensic software empowering police departments in more than 60 countries. yug@belkasoft.com*

Oleg Afonin *is an expert in digital forensics, a researcher in the area of digital security, and a specialist in data recovery. research@belkasoft.com*

Source URL (retrieved on *11/23/2014 - 6:14pm*):
http://www.dfinews.com/articles/2014/10/recovering-evidence-ssd-drives-understanding-trim-garbage-collection-and-exclusions

Links:
[1] http://www.dfinews.com/articles/2012/09/why-ssd-drives-destroy-court-evidence-and-what-can-be-done-about-it-part-1
[2] http://en.wikipedia.org/wiki/Trim_(computing)
[3] http://www.overclock.net/t/1330730/ocz-firmware-2-25-trim-doesnt-work-bug-regression-bad-ocz-experience
[4] http://www.ocztechnologyforum.com/forum/showthread.php?96382-Deterministic-Read-After-Trim
[5] http://www.extremetech.com/extreme/184253-ssd-shadiness-kingston-and-pny-caught-bait-and-switching-cheaper-components-after-good-reviews
[6] http://www.anandtech.com/show/6477/trim-raid0-ssd-arrays-work-with-intel-6series-motherboards-too