Fundamentals
of
Storage Media Sanitation

Tom Olzak
June 2006

One of the most fundamental principles of information security is that it's all about the data. Data in transit or at rest is the primary focus of administrative, physical, and technical safeguards. Security professionals are doing better every day when it comes to protecting information in static production environments. But what happens when magnetic, optical, or semiconductor media is repurposed or retired?

In this paper, I define media sanitation and how it fits into an overall security program. Next, I examine how attackers can extract information from electronic media—even after it's been overwritten. Finally, I explore ways you can protect your organization from attacks—both casual and highly motivated.

What is Media Sanitation?

When electronic media is repurposed or retired, it's the responsible of the data owner—the representative of the responsible organization—to ensure the data currently or previously stored on that media is not easily accessible. This is the purpose of media sanitation. Put another way, media sanitation

"…refers to the general process of removing data from storage media, such that there is reasonable assurance, in proportion to the confidentiality of the data, that the data may not be retrieved and reconstructed" (Scholl et al, 2006).

Media sanitation is another control, or safeguard, that should be implemented in accordance with risk management principles. Reasonable assurance varies based on the sensitivity of the information and the retrieval methods most likely to be used by an attacker. The method used depends on the reward or value derived from obtaining the data versus the work factor and costs associated with available retrieval resources.

There are two basic categories of data retrieval—keyboard and lab. The keyboard category includes all normal data access methods. An example of a normal access method is the use of software installed on the system housing the information, whether authorized or unauthorized, while located in it normal-use location. The lab category includes sophisticated retrieval techniques, both hardware and software, that usually require the physical removal of the storage media from the secure location where it's normally used. The retrieval resources employed are installed and configured in an off-site lab environment.

Each media type might have its own unique set of sanitation challenges. I chose three types to use as examples as we work through the various difficulties encountered when attempting to protect data on a data storage medium in transition. These types include magnetic disks, optical disks, and semiconductor devices.

Before looking at specific storage media types we need to understand the underlying challenges common to all media—file deletion and data remanence.

File Deletion

Deleting files from storage media doesn't necessarily delete the data contained in the file. Instead, the file names—or some other pointers to the files—are marked as deleted. This tells the controlling operating system (OS) that it can reuse the areas associated with the deleted files. Until a deleted file's data is overwritten, it's still subject to retrieval by keyboard methods with off-the-shelf or OS-based utilities.

A single overwrite of data is a good way to prevent keyboard retrieval of sensitive information. But lab retrieval defense requires more. This is due to data remanence.

Data Remanence

"Data remanence is the residual physical representation of data that has been in some way erased" (NCSC, 1991). In other words, data that has been removed or overwritten once or twice is potentially retrievable through the use of lab-based methods. Data remanence has various causes.

Magnetic Disks

Figure 1 is a photo of the internal components of a 5.25" hard drive. The topmost disk (platter) is on the left. On the right are the components that make up the disk arm assembly. Note the disk arm hovering over the platter. At the end of the disk arm is the read/write head.

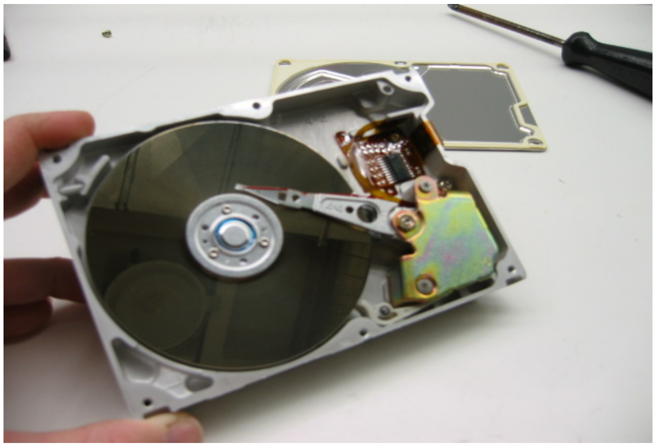

Figure 1: Hard Drive Internals (Niagara College, 2006)

A typical hard drive consists of multiple platters connected to a central spindle (Figure 2). Each platter surface is logically divided into tracks. Tracks are further divided into sectors. A read/write head floats above each platter surface. When a file is written, the drive selects sectors on one or more available tracks on one or more platters and writes the file data. The information about where the file was written is stored in a directory located on one of the platters. (This is a very simplified explanation of how hard drives work. For more detailed information, see the Wirzenius, Oja, and Stafford reference in *Works Cited*.)

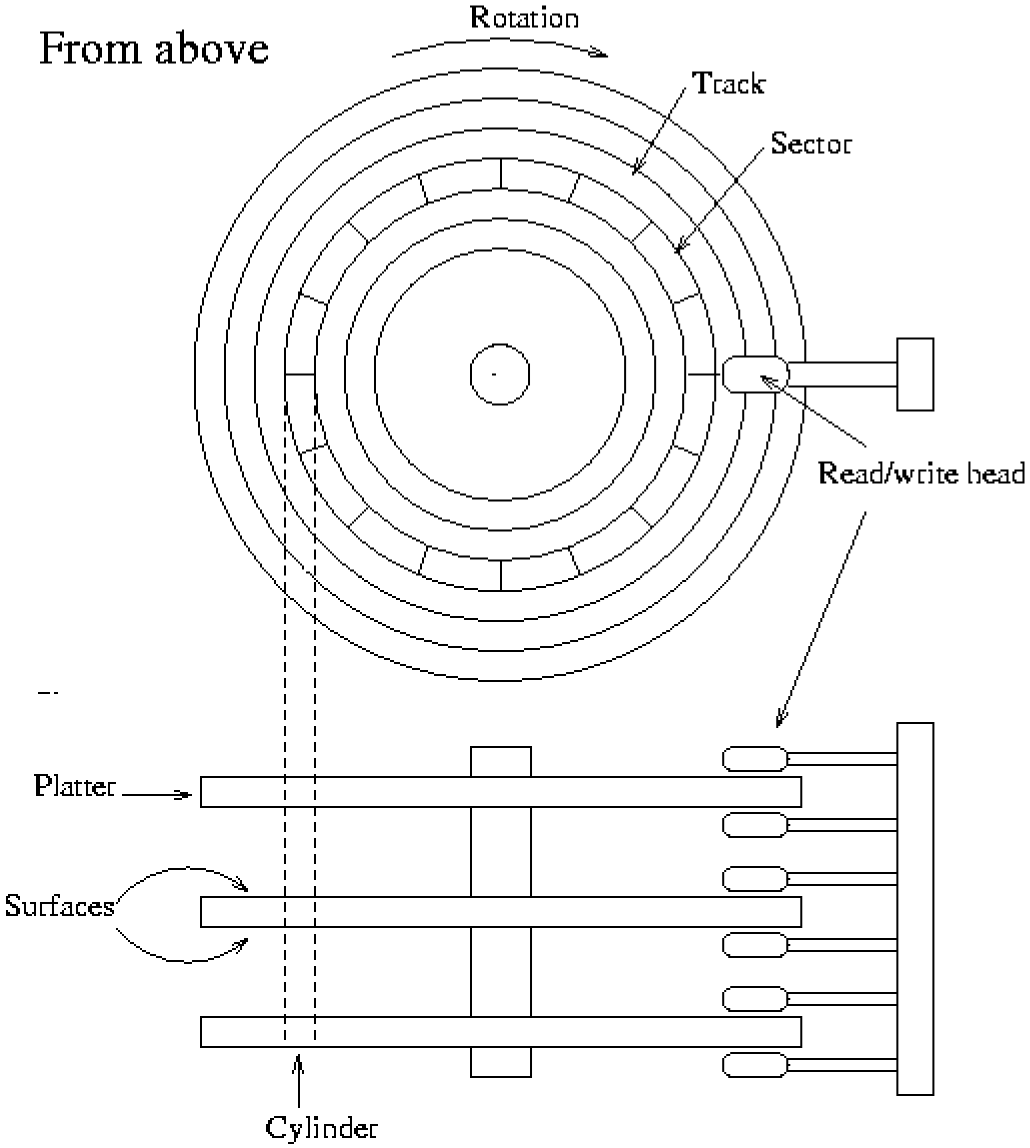

Figure 2: Hard Drive Schematic
(Wirzenius, Oja, and Stafford, 1993-2001)

When a file is deleted, the information about how to locate the file is marked as deleted in the directory. However, the file data still resides in the tracks and sectors on the platter(s). When future files are written to the drive, one or more of the sectors occupied by the deleted file might be overwritten with new data, but until all the sectors are overwritten any attacker can use common tools to retrieve part or all of the "deleted" information. Retrieval of deleted files in this way is a keyboard method. No special lab equipment is required.

One way to eliminate the possibility of successful keyboard attacks is to overwrite the deleted file's disk sectors. However, this isn't a perfect solution. After a single overwrite, lab retrieval technology can "read" the data previously written to the disk. This is due to variations in the strength of the recorded bits as well as stray magnetism at the edges of the tracks (Gutmann, 1996). The number of overwrites necessary depends on the sensitivity of the information you're trying to protect. According to Gutmann,

"When all the… factors are combined it turns out that each track contains an image of everything ever written to it, but the contribution from each "layer" gets progressively smaller the further back it was made. Intelligence organisations have a lot of experience in recovering these palimpsestuous images" (1996)

Advances in disk technology are making it harder to recover overwritten data, but it's still an issue for magnetic disks--containing highly sensitive information--that are removed from the protection of a secure physical environment.

One final note about overwriting data stored on magnetic disks. As a drive ages, sectors once used for data storage fail to meet working parameters as seen by the drive electronics. These sectors are marked as bad. During overwrite processes, these sectors are normally ignored, leaving the data stored there available for potential retrieval.

Optical Disks

All optical disk technologies (CD-ROM, write-once, rewritable, etc.) work in essentially the same way. A read laser detects light and dark areas on the disk surface. The light (reflective) areas cause the low intensity beam from the laser to reflect back into a read head. These pulsing reflections are translated into 1's and 0's.

The only effective way to dispose of non-rewritable optical media is destruction. Data on rewritable optical disk can be overwritten. However, there isn't enough empirical evidence to support the premise that overwriting optical media is effective in preventing lab attacks. Like other types of optical disk media, destruction appears to be the best course of action when a CD holds highly sensitive information.

Memory

We usually consider Random Access Memory (RAM) as an easy medium to erase. After all, all you have to do is remove power and everything stored immediately disappears. And erasing flash memory is as easy as deleting all data. What could be left behind? The answer is, it depends.

According to Gutmann, semiconductor devices can "remember". In other words, the characteristics of the substances that make up memory components create a tendency in those components to retain a trace of bits previously stored. The strength of this trace retention depends on the length of time the data occupied the same memory location (2001).

In RAM, sensitive pieces of data might be stored for long periods. For example, encryption variables are often stored in the same place in memory for as long as the system is powered on. This might cause a trace of the crypto information to remain after power is removed. Even if the memory is reused, data previously stored might be retrievable. The memory that makes up thumb drives has a similar problem. Simply erasing or attempting to overwrite data stored on your USB thumb drive might not be sufficient to protect highly sensitive information from lab retrieval systems.

Guidelines for Media Sanitation

According to the NIST Guidelines for Media Sanitation, there are four sanitization processes (Scholl et al, 2006)—disposal, clearing, purging, and destruction.

- **Disposal –** The process of disposal essentially consists of tossing the media in a dumpster with no attempt to hinder or prevent the recovery of data. This also includes the reuse of disks, tape, or memory without taking steps to protect information that may have been stored during previous operational use. It's acceptable to follow a simple disposal process when the data stored on the media is classified as "public". In other words, the release of the information will not cause harm to the organization, its employees, its shareholders, or its customers.

- **Clearing –** Clearing requires taking steps to prevent the recovery of data through a keyboard attack. As we've seen, this requires more than deleting files. At least a single overwrite of the writable areas of the media must be completed. A single overwrite significantly increases the effort, or work factor, required to recover information. Not only does the attacker need physical access to the media, but recovery requires the use of lab-based tools. Clearing is acceptable when release of the information stored would cause only moderate harm to the organization, its employees, its shareholders, or its customers.

- **Purging –** Purging is necessary when the compromise of the information stored on the media will result in serious--and possibly irrecoverable--harm to the organization, its employees, its shareholders, or its customers. Data is overwritten enough times to increase the work factor of lab-attack attempts to a level that exceeds the data's value to the attacker. Ideally, all remanent data is removed.

- **Destroying –** Purging is a good way to retain media you wish to reuse. However, the best process for ensuring the irretrievability of highly sensitive data is to destroy the media. During the destruction process, media is reduced to a state in which both keyboard and lab attack attempts are impossible.

Again, the process you select depends on the sensitivity of your information and the potential impact on your business if the information is compromised. With these basic processes in mind, let’s look at specific approaches to sanitizing magnetic, optical, and semiconductor storage.

Magnetic Media

Clearing magnetic storage is a simple matter of writing a single character to all writable areas of a disk or tape. This prevents the use of easily obtainable utilities to recover deleted files. Purging is not so easy.

Purging requires that all usable remnants of any data ever stored on the tape or disk is irretrievable. Earlier in this paper, we looked at Gutmann’s assertion that successful data retrieval grows less probable the more times it’s overwritten. Further, recovering information is made possible by calculating variances in voltage levels detected by the read head.

Effective purging, using an overwrite technique, consists of two factors. First, the data must be overwritten a sufficient number of times to make recovery very difficult, if not impossible. Second, the overwrite cycles must use alternating 1’s and 0’s. For example, if we wrote all 0’s to all writable areas on a tape during the first overwrite pass, we would write all 1’s on the second pass. However, there is a problem with this approach.

Certain disk technologies will not write a large number of contiguous 0’s or 1’s. Why is outside the scope of this paper. What is important to understand is that the tool you use must take this into account. The best approach is alternating bit sequences that result in writing the complement of the bit written during the previous write cycle. The following table lists possible bit sequences that meet the necessary criteria for magnetic storage purging.

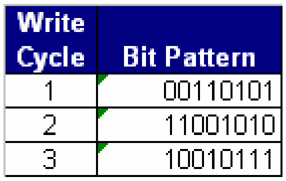

Write Cycle	Bit Pattern
1	00110101
2	11001010
3	10010111

This series of patterns meets the Department of Defense general standard for purging magnetic media which is:

1. write a single pattern
2. write its complement
3. write another pattern

The actual number of these overwrite cycles necessary for tape or disk depends on the storage media and its sensitivity (NCSC, 1991).

Degaussing is another way to purge magnetic media by erasing all data ever stored on a tape or disk. Degaussers use a electromagnetic field to destroy magnetic imprints on media. The degausser used depends on the media processed. Various magnetic field strength levels are required, depending on the media types. Also, degaussers must be serviced regularly to ensure they continue to produce the expected field strength. Degaussing isn't always the best approach when you want to reuse the erased media. Exposing certain types of tapes and hard drives to a strong magnetic field will render them useless. Be sure to check with the media manufacturer.

Finally, magnetic media you don't plan to reuse can be destroyed. Acceptable destruction methods include pulverizing, smelting, incineration, and shredding.

Optical Disks

If clearing is your objective, overwriting re-writable optical disks might be acceptable. But again, there is no proven method for purging them. Further, overwriting is impossible for clearing or purging other types of optical media. When dealing with highly sensitive information stored on optical disks, destruction is your best option.

In addition to the destruction options listed above for magnetic media, you might also apply an abrasive substance (i.e., an emery wheel or sander) to the recording surface. There are products available that allow you to feed stacks of optical disks into a device that makes this approach quick and relatively easy.

Memory

Preventing semiconductor data remanence begins before the storage media is ever used. Gutmann recommends the following steps to reduce the potential threats posed by semiconductor data remanence (2001).

- Don't store cryto-keys in the same RAM location for long periods. Occasionally move them to different locations and clear the original location.
- Cycle EEPROM/flash cells 10 to 100 times with random data before writing anything sensitive to them to eliminate any noticeable remanence effects arising from the use of fresh cells.
- Don't assume that a key held in RAM in a piece of crypto hardware is destroyed when the RAM is cleared. The circuitry might carry an after-image of the key.
- Remember that some non-volatile memory devices are a little too intelligent, and may leave copies of sensitive data in mapped-out memory blocks after the active copy is erased.

Overwriting all memory cells is an acceptable method to clear semiconductor memory. However, destruction or degaussing might be the only processes your organization find acceptable for purging. Again, it depends on the media and the purging tools used. The destruction techniques listed for magnetic media also apply to memory devices.

Conclusion

Storage media manufacturers are not ignoring data remanence challenges. Each generation of storage media includes magnetic and semiconductor devices that are better at resisting remanence than the previous generation. The technology exists today that increases lab-attack work factors to significant levels, but it typically resides on "…the newest, highest-density (and by extension most exotic) storage devices available" (Gutmann, 2001).

Like most of you, I work in the real world. Trying to make a business case for purchasing sophisticated devices or to pay a third party for media destruction isn't always easy. Like any other information security proposal, the solution must be based on a solid risk management assessment. If the potential impact on the business doesn't justify the costs associated with media sanitation, then it's probably not the right thing to do at this time. However, it's a good idea to ask the right questions as your network engineers start ripping out those old servers…

Tom Olzak, MBA, CISSP, MCSE, is President and CEO of Erudio Security, LLC. He can be reached at tom.olzak@erudiosecurity.com. Additional security management resources are available at http://adventuresinsecurity.com.

Check out Tom's book, *Just Enough Security*, at Amazon.com

Visit Tom's blog at http://blogs.ittoolbox.com/security/adventures/

Listen to Tom's podcasts at http://adventuresinsecurity.com/podcasts

Free security training available at http://adventuresinsecurity.com/SCourses

Works Cited

Gutmann, P. (1996). *Secure deletion of data from magnetic and solid–state memory.* Retrieved May 15, 2006 from http://www.usenix.org/publications/library/proceedings/sec96/full_papers/gutmann/index.html

Gutmann, P. (2001). *Data remanence in semiconductor devices.* Retrieved May 15, 2006 from http://www.usenix.org/events/sec01/full_papers/gutmann/gutmann.pdf

NCSC (1991). *A guide to understanding data remanence in automated information systems.* Retrieved May 12, 2006 from http://www.radium.ncsc.mil/tpep/library/rainbow/NCSC-TG-025.2.txt

Niagara College (2006). *Future trends in technology (TECH1271), unit 5 course notes.* Retrieved May 12, 2006 from http://www.technology.niagarac.on.ca/courses/tech1271/Unit5.html

Scholl, M., Kissel, R., Skolochenko, S., & Li, X. (2006, February). *Guidelines for media sanitation (NIST SP 800-88, Public Draft).* Retrieved May 20, 2006 from http://csrc.nist.gov/publications/drafts/DRAFT-sp800-88-Feb3_2006.pdf

Wirzenius, L., Oja, J., & Stafford, S. (1993-2001). *The Linux administrator's guide chapter 6.* Retrieved May 12, 2006 from http://www.faqs.org/docs/linux_admin/x1001.html