The Difficulty of Data Annihilation from Disk Drives: or Exnihilation Made Easy

Dean Devera
ddevera@cox.net
CS 574 / Tom Perrine
11 December 2001

> "At length my eyes, in going the circuit of the room, fell upon a trumpery filigree card-rack of pasteboard, that hung dangling by a dirty blue ribbon, from a little brass knob just beneath the middle of the mantelpiece. In this rack, which had three or four compartments, were five or six visiting cards and a solitary letter . . . No sooner had I glanced at this letter, than I concluded it to be that of which I was in search." (Poe 1844)

In Edgar Allen Poe's "The Purloined Letter", a building is literally torn apart in an attempt to find a note of immense political consequence. What the desperate searchers failed to realize was that the entire time, the letter was shoved in a card-rack . . . hidden in plain sight. Unrelated to deep analysis of disk drives? Not quite. Unlike Poe's baffled searchers, computer forensics experts examining hard drives oftentimes know the evidence they seek resides right beneath their noses, stored in magnetic patterns on the drive's platters. But to bring this evidence to light may require as much effort as expended by the seekers of the purloined letter.

Initially, this paper will examine some of the more esoteric methods employed for recovering data from inert disk drives. This portion will be primarily expository, with the weight of analysis being done toward the end of the paper, where the present effectiveness of these methods in light of recent

advancements in hard drive technology is assessed. Regarding the latter, the paper will attempt to answer the following question: Is there anything one can do to assure “erased” data is completely, irrevocably *annihilated*?

The name often associated with this sort of deep analysis is Peter Gutmann. In 1996, Gutmann synthesized the then-existent relevant knowledge concerning magnetic media analysis with his own, publishing the seminal paper, “Secure Deletion of Data from Magnetic and Solid-State Memory”.

In it, he prescribes high-powered magnetic microscopy for probing and imaging the magnetization patterns found on the surface of disk media. The two techniques he describes, magnetic force microscopy (MFM) and magnetic force scanning tunneling microscopy (STM), afford the user a relatively simple means for extracting images off the surface of a drive platter.

How simple? I quote Gutmann in full, and remind the reader that his estimate is based on technology that’s already had half a decade to grow. “Even for a relatively inexperienced user the time to start getting images of the data on a drive platter is about 5 minutes. To start getting useful images of a particular track requires more than a passing knowledge of disk formats, but these are well-documented, and once the correct location on the platter is found a single image would take approximately 2-10 minutes depending on the skill of the operator and the resolution required.” (Gutmann 1996)

There are certain inherent characteristics of hard drive design, specifically track layout, which admit this kind of analysis. One concerns the width of the tracks in comparison to the “width” of the actual data written to them. Put another way, if the read/write heads paint 1s and 0s in a swath as wide as your typical freeway lane, the track itself would be equivalent to a three-lane highway (Preston 1995).

What this means is the vestiges of “overwritten” data may persist toward the edges of the data track. These can, in turn, be detected and measured via methods employing MFM and/or STM. In fact, these same techniques are used by drive manufacturers themselves to verify if a unit’s servos are functioning correctly (Gomez 1992).

Even if the vestiges are indistinct, recovery isn’t impossible. Gutmann writes that “the new track width can exhibit modulation which depends on the phase relationship between the old and new patterns, allowing the previous data to be recovered even if the old data patterns themselves are no longer distinct.” (Gutmann 1996)

Given the rapid increase in hard drive capacities since the ’96 paper (a 100 gigabyte unit is available at the time of writing) one may question the present effectiveness of this method. As platter densities are upped to increase the maximum capacity of the drive, the width of the tracks on a fixed-sized platter must shrink, meaning there’s less space on the “shoulders” of each track for digital debris from previous writes to accumulate. The effect of this would presumably be more pronounced in the highest capacity drives, making them more resistant to this type of analysis. In a follow-up paper, Gutmann implies this may be so. (Gutmann 2001)

The previous does nothing to dilute the paper’s most astonishing claim. Gutmann writes, “When all the above factors are combined it turns out that each track contains an image of everything ever written to it, but that the contribution from each ‘layer’ gets progressively smaller the further back it was made” (Gutmann 1996). Regarding overwrites, perhaps the high capacity drives of today only afford a shorter “history” than those of 5 years ago, and nothing more in the way of complete and irrevocable erasure of old data by new.

All this doesn’t bode well for someone intent on committing sensitive information to disk. Most are familiar with software-based solutions that claim to erase data (e.g. Evidence Eliminator and the like). But what software can avert future analysis by a magnetic microscope? Granted, the assumption is the average person needn’t worry about the disks being ripped out of their system for forensic analysis. But such concern ought to be commensurate with the sensitivity of the information stored. Can one be assured that highly sensitive data can be completely eliminated should a situation warrant it? Put another way, is there anything one can do to thwart Gutmann's bag of tricks? Anything *plausible*? Maybe. And it’s the very progress of drive technology which may abet the miscreant.

In the absence of any specialized equipment, the most one can do is follow the complex overwriting procedures described by Gutmann in the ’96 paper. In light of the fact that vestiges from previous writes may persist, these ought to be followed to the letter (Preston 1995). However, as stated before, the increase in drive capacity and the subsequent narrowing of track width may make these procedures more than sufficient . . . perhaps even overkill. That said, repeated overwrites do not offer as much assurance as, say, actual physical destruction of the drive platters.

Let's examine the hypothetical case of a miscreant who, well aware of the previously mentioned analysis techniques, takes a hammer to the hard drive. Better yet, let's assume he's removed the hard drive’s aluminum casing and is swinging away at the platters themselves. When reconstructing broken platters, it seems the most difficult task would be properly realigning those micron-width tracks. Piecing together said platters would be akin to working on the world's toughest jigsaw puzzle.

If the platters were only fractured radially, the task becomes a bit easier. But realignment remains a nightmare, and could spell the difference between a drive that can offer up evidence to the forensic specialists and a drive chock full of digital gibberish. It's exceedingly difficult, but not impossible if we're dealing with relatively few pieces.

Once reassembled, high-powered magnetic microscopy could then be turned to the media surface. However, the evidential weight of information gleaned from such a precarious operation remains unclear. After all, the specialist is constructing a puzzle that lacks a nice box-top depiction of what the end product is supposed to resemble. The specialist has no idea what the digital sentences laid out on each track are supposed to read, lending another difficulty to track realignment. In addition, issues of evidence tampering creep in. Most suspect drives are imaged before any analysis takes place, with the original stored in a locked safe to skirt such issues (Perrine 2001).

But the possibility of platter reconstruction exists, however minute. And piecing together a shattered disk may not be such an eccentric undertaking if the data we're trying to get at is of extreme consequence (e.g. top-secret information, either corporate of governmental). So destroying a drive's guts with a hammer fails in providing complete assurance that data is completely annihilated.

What about degaussing the platters with a magnetic source many times more powerful than the drive's erase head? Given evidence supporting the persistence of "erased" data, the source would have to be quite powerful. It's been demonstrated that a magnetic force five times greater than the "coercivity" of the medium is required to erase data beyond all realistic hope of recovery (Preston 1995). "Coercivity is the force necessary to magnetize each tiny data

area and to record a 1 or 0 there. The higher the coercivity, which is measured in Oersteds (0e), the harder the data is to erase." (Preston 1995)

Gutmann adds that, "Since degaussers tend to be rated by whether they erase sufficiently for clean rerecording rather than whether they make the information impossible to recover, it may be necessary to resort to physical destruction of the media to completely sanitise it." (Gutmann 1996)

It seems the equipment required for data destruction via degaussing is a bit much . . . perhaps too much for the average person to provide, and economically prohibitive to hire someone else who has the right equipment to do it. One such example of an "adequate" degausser, circa 1991, was a 2.5-megawatt Navy research magnet that not only succeeded in complete erasure. It *bent the drive platters* as well (Hayes 1991). So they aren't things the average consumer can readily purchase. Simple DC erasure also runs into the same wall thrown up by the inherent coercivity property of the media (Adly 1993).

What's left is torching the drive. Obviously, electron microscopy proves fruitless when trying to extract 1s and 0s from a smoldering lump of metal. And even Gutmann recommends arson when it comes to floppy discs (Gutmann 1996).

I ask the reader's indulgence if the previous ruminations seem flippant. But I'm trying to emphasize an important point. Gutmann has shown that short of melting the platters, complete and irrevocable erasure of data is close to impossible. The implications of this are manifest. Computer forensic specialists will always have a place to search for evidence. In other words, our hypothetical miscreant can never be sure his tracks were completely swept away.

However, the aforementioned advancements in hard drive capacity do offer a sliver of hope. As drive capacities and, subsequently, densities increase, so does the difficulty in extracting information via microscopy become, whether the

platters are shattered or fully intact. Gutmann says so himself in a follow-up to the ’96 paper: “Finally, however, the best defense against data remanence in semiconductor memory is, as with the related problem of data stored on magnetic media, the fact that ever-shrinking device dimensions, and the use of novel techniques such as multilevel storage is making it more and more difficult to recover data from devices.” (Gutmann 2001)

There’s an alternative spin one can take when thinking about the difficulty of data destruction. Given that the relevant government agencies are aware of the exotic techniques described earlier (maybe even before Gutmann spelled them out . . . recall the “discovery” of differential cryptanalysis in the early ‘90s!), it’s interesting to speculate how the “good guys” go about destroying data. For obvious reasons, no well-documented material describing their methods exists for public consumption. However, it stands to reason that protocol must dictate, at the very least, strong encryption of sensitive data before it’s written to disc, if it’s even allowed to be committed to disk at all.

Granted, these techniques are highly specialized and time-consuming (read: expensive). But they grow more warranted with the seriousness of the alleged offense. If evidence extraction from hard drives is a matter of national security, the aforementioned restraints become irrelevant.

Yet there are those that will continue to thumb their noses at Gutmann's claims of digital *exnihilation*, either out of ignorance, bravado, or a combination of both. To those I have a single recommendation: get a blowtorch.

References

Adly, A., Mayergoyz, I., Burke, E. (November 1993). "Computation of Magnetic Fields in Hysteretic Media". *IEEE Transactions on Magnetics, Vol.30, No.6*. p.4248

Gomez, R., Adly, A., Mayergoyz, I., Burke, E. (September 1992). "Magnetic Force Scanning Tunneling Microscope Imaging of Overwritten Data". *IEEE Transactions on Magnetics, Vol.28, No.56*. p.3141

Gutmann, P. (July 1996). "Secure Deletion of Data from Magnetic and Solid-State Memory". *Sixth USENIX Security Symposium Proceedings, San Jose, CA, July 22-25, 1996*.

Gutmann, P. (2001). "Data Remanence in Semiconductor Devices". *IBM T.J. Watson Research Center*. http://www.cryptoapps.com/~peter/usenix01.pdf

Hayes, D. (July 1991). "How many times erased does DoD want?". posting to comp.periphs.scsi newsgroup. message-ID 1991Jul24.050701.16005@sulaco.lonestar.org.

Perrine, T. (November 2001). "Intrusion Analysis (Compromised Linux Box)." Computer Science 574. San Diego State University. November 29, 2001.

Poe, E.A., (1845). "The Purloined Letter". *Tales*. pp. 200-218. http://www.eapoe.org/works/tales/plttrb.htm

Preston, C. (February 1995). "The Data Dilemma". *Security Management Journal*. www.usarc.army.mil/dcsint2/documents/security/Security%20Regs/Information%20Systems%20Security/Datadilemma.doc