CMRR Protocols for Disk Drive Secure Erase

Users want and need a simple and secure way to erase all their data from disk drives, when releasing them from their physical control for resale or repair. Over a third of drives resold on eBay contain personal data such as credit and medical records.[1] These drives come from PCs, servers, ATM machines, banks, and workstations.

It is important to initially emphasize that erasure security can only be relative. There is no method giving absolutely secure erase. Government document DoD 522.22M is commonly quoted on erasure methods, and requires physical destruction of the storage medium (the magnetic disks) for data classified higher than Secret. However, even such physical destruction is not absolute if any remaining disk pieces are larger than a single 512-byte record block in size, about 1/125" today's drives. Pieces of this size are found in bags of destroyed disk pieces studied at CMRR. Magnetic microscopy can image the stored recorded media bits, using the CMRR scanning magnetoresistive microscope. Physical destruction nevertheless offers the highest level of erasure because recovering any actual user data from a magnetic image requires overcoming almost a dozen independent recording technology hurdles. This is an example of "exotic time consuming technology" necessary as the barrier to data recovery for the highest level of erasure security. Even if these hurdles were overcome, about an hour would be required to recover one single user data block out of millions on the disk. Recovering substantial amounts of data in less than months requires that the disk be intact and undamaged so that heads can be flown over it to obtain data playback signals, and also overcoming the technology hurdles. Simply bending a disk makes this impossible.

CMRR has established and tested protocols for software secure erase. Their security levels vary between the levels just discussed. Four basic security levels are defined, Weak erase, block erase, Normal secure erase, and Enhanced secure erase. Block and Normal secure erase are intended for elimination of user data up to the Secret level, and Enhanced secure erase for higher levels. The Enhanced level is not yet in drives for verification.

These four erasure protocols exist is because users make a tradeoff between the erasure security level and the erasure time required. A high security protocol requiring custom software or days to accomplish will be avoided by most users, making it little used and therefore of limited practical value. For example, DoD 5220 calls for multiple block overwrites for Secret data, which can take more than a day to complete in today's drives. So users make a tradeoff between the time required to eliminate his data and the risk that the next drive user will know and use recovery techniques to access weakly erased data. Figure 1 shows tradeoffs in security level vs. speed of erasure for various erasure options.

---

[1] S. Garfinkel, A. Shelat, "A Study of Disk Sanitization Practices," IEEE Security and Privacy, January-February 2003

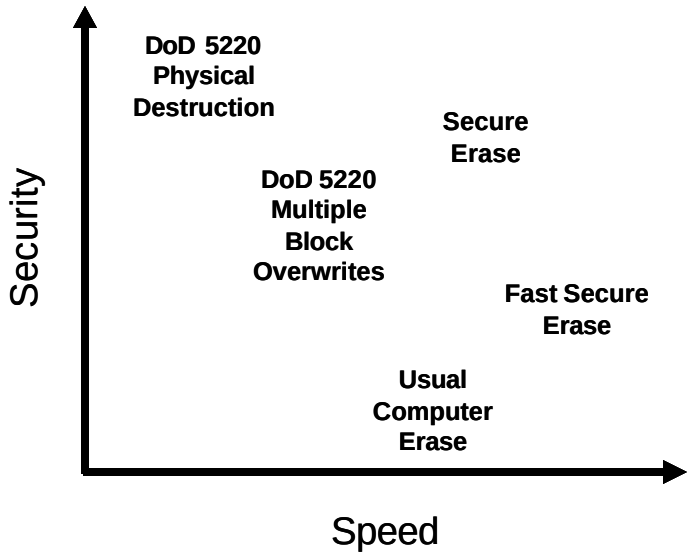

Figure 1. Security vs. speed of completion of various modes to erase data on hard disk drives.

For all but top-secret information and when time is critical, users will often turn to erasure that takes minutes rather than hours or days. They will select a method giving them an acceptable level of security in a reasonable time window.

CMRR has studied secure erase for the Federal Government for seven years, and its research[2] shows three distinct protocols used for user data deletion, with the following security levels:

(1) **Weak deletion** by users deleting files in public operating systems such as Windows or Linux ("usual computer erase' in Figure 1). This deletes only file directory entries, not the user data itself. User data recovery is simple with utilities like Norton Unerase. Even reformatting a drive still only erases file directories, not user data. Low level formatting commands no longer exist in today's disk drives.

(2) **Block erasure utilities** are widely available for purchase, which overwrite all user accessible blocks. Block overwriting gives a higher level of deletion confidence than (1) and these utilities claim to meet Federal Government requirements in DoD 5220. This document requires three writes – 0's, its binary complement 1's, and a then random data pattern which is verified by a read. However, block write software utilities cannot erase reassigned user blocks, since these have no logical block address to write to and physical sector address drive commands no longer exist. Some utilities do not verify the final random write; and all can be vulnerable to malicious software attacks which modify the utility program to make it falsely report a successful erase. Operating system erasure commands using internal kernel security could eliminate false reporting risk, but MS Windows does not currently have such a secure erase commands. In Linux, a simple shell sequence can be run to overwrite all accessible disk blocks, but its security is not established.

(3) **Disk drive Secure Erase** is a drive command defined in the ANSI ATA and SCSI disk drive interface specifications, which runs inside drive hardware. It completes in about 1/8 the time of 5220 block erasure. It was added to the ATA specification in part at CMRR request. All recent ATA drives have the command

[2] G. Hughes, T. Coughlin, "Secure Erase of Disk Drive Data" IDEMA Insight Magazine, Spring 2002

(>10-15 GB) and successfully pass secure erase validation testing at CMRR (see Appendix). The next section covers its technical requirements for erasure security. The command reports whether the secure erase is totally successful, through the ATA hardware interface. It has DoD 5220 Secret data erasure security and offers an opportunity for higher erasure security, if the Enhanced protocol requirements below are met. Its erase and malicious attack security meets DoD 5220, because attacking drive internal firmware is far more difficult than attacking computer software, and requires disk drive forensic technology. A standardized internal secure erase command also exists for SCSI drives, but it is optional and not currently implemented in SCSI drives tested by CMRR. SCSI drives are a small percentage of the world's hard disk drives, and the command will be implemented when users demand it.

Secure Erase Requirements

Secure Erase must be a command to a standardized interface data storage device, defined by its ANSI specification (American National Standards Institute). The device must internally perform the erase and report successful erase only if all user accessible records have been erased. Reassigned blocks may optionally not be overwritten, if no subsequent user access to them is possible. The command should allow security against a malicious attacker erasing a drive.

Normal Secure Erase

Secure erase is called Security Erase in the ATA disk specification and Security Initialize in SCSI. The command must cause an overwrite operation that stores random bits in all user accessible blocks on storage media. The overwriting user data itself need not be random if the device randomizes user bits before media storage.

The current ATA specification for Normal Erase mode states that the SECURITY ERASE UNIT command shall write binary zeroes to all user accessible data areas. (ATA reassigned blocks are not user accessible because they have no user address). This level of erasure is excellent for fast erasure, although it does not precisely follow the three writes called out in DoD 5220. CMRR verification testing (below) showed that the erasure security is at the level of DoD 5220, because drives having the command also randomize user bits before storing on magnetic media. In-drive block verify is via internal write fault detection hardware, which takes no additional time thus increases user willingness to use the command. The three block writes of DoD 5220 plus verify can take far longer than the secure erase command. CMRR test times were up to days but the drive normal Secure Erase can complete in 30-45 minutes.

Enhanced Secure Erase Requirements

The current ATA ANSI specification states: *"When Enhanced Erase mode is specified, the device shall write predetermined data patterns to all user data areas. In Enhanced Erase mode, all previously written user data shall be overwritten, including sectors that are no longer in use due to reallocation. This command shall disable the device Lock mode, however, the Master password shall still be stored internally within the device and may be reactivated later when a new User password is set.*

CMRR has established minimum mandatory properties of an Enhanced Secure Erasure algorithm which provide erasure security equivalent to most implementations of physical destruction and in a much shorter time. CMRR specifies a minimum of two random data writes of all physical user sectors (including reassigned sectors), where each write is offset off-track opposite to the other by at least 10% of the track pitch. The number of physical sectors not successfully written should be reported, and any defective sectors which could not be

written. Reassigned blocks may optionally not be overwritten, if no subsequent user access to them is possible. (Example: a block reassigned because of a hard error in its embedded servo, causing a write fault which prevents writing the data block).

Overwritten data left in track edges is normally unreadable magnetic noise[2], but the offtrack writes makes any possible coherent data in the track edges unrecoverable. Note that only drive internal technology is able to accomplish an offtrack Secure Erase. There is no standardized "write offtrack" command for any software utility to use.

A proposal for validation of Enhanced Secure Erase is described in the Appendix, and would additionally include examination of disks in the CMRR scanning magnetoresistive microscope. No image traces of magnetic bits must be found, including track edges.

APPENDIX: Secure Erase Validation for ATA disk drives

ATA Security Erase:

Secure Erase is a mandatory part of the password protection command set in the ATA disk drive specifications. (The overall password command set is optional, but universal in commercial ATA drives.) Security Erase can be executed in either normal or enhanced mode. In normal mode, the device must erase all user accessible data areas by overwriting with data zeros. In enhanced mode the device must overwrite all user accessible data areas with a predetermined data pattern, and in addition overwrite all reallocated sectors. A Security Erase command can be issued by either device user or master, using their respective drive UnLock passwords. The user/master passwords are set by the security Set Password command, which also sets the security level (high/maximum).

Security Erase validation testing an ATA device determines whether:

- All user data is erased, beyond recovery through the ATA interface
- Reallocated block data is irrecoverable
- The Security Erase performed was single or multipass, low frequency or random data patterns.
- The device supports *Fast* Security Erase (data is inaccessible until a new user completes Security Erase).

These command validation tests do not establish whether exotic recovery of erased data is possible. That requires additional testing using scientific magnetic microscopy instrumentation on disassembled drive disks.

Validation Test Protocol:

1. Device configuration test:
   Commands are issued to the ATA drive under test to check correct device function including soft reset, read and write commands.
   - A Device Configuration Identify command is then issued, which returns whether the device supports the 48-bit command set and the security mode features. If the device doesn't support the ATA security command set,
   - A Device Configuration Set command is issued to enable the security feature set. If this fails, then testing is terminated with status of Security Erase failure.

2. Security Erase command functionality:
   Device identify commands are issued, determining:
   - If the device supports enhanced Security Erase
   - The time required for normal Security Erase completion
   - The time required for enhanced Security Erase completion.
   - Whether Security Erase ends with device security enabled and whether the device is in locked state or unlocked state. (In the locked state read or write commands are rejected by the device, until a security unlock command is executed with correct password)
   - The current security level (disabled, high, or maximum)
   - Whether the device is in frozen state. In the frozen state all security related commands are rejected.

3. Normal Security Erase test:
   Normal Security Erase command is tested for 8 cases:
   - User password set for high security level
   - User password set for maximum security level
   - Master password set for high security level
   - Master password set for maximum security level.

   For each test, all user accessible drive sectors are filled with a unique test pattern and verified. A Security Erase command is then issued by user and master on each of the above combinations. The data pattern written on the drive is checked using read sector commands, looking for any unerased unique data patterns. The time required for the drive to do secure erase is logged.

4. Enhanced Security Erase test:
   If enhanced Security Erase is supported, Security Erase is issued for each of the above 8 cases in enhanced mode, the data pattern written on the drive is retrieved, and the time for the enhanced Security Erase is logged.

5. Interrupting Security Erase:
   Security Erase procedure is interrupted during execution by removing device power. Device state after repowering shall be “drive locked”

6. Fast Erase
   After execution, verify that drive reboots in Locked state and rejects data access commands.

Validation Test Hardware and Software Setup:

An Intel P4 2 GHz computer is used at CMRR. For each test, scripts send appropriate commands and to process the data returned. The scripts were written using the ATA device driver code ATADRVR written by Hale Landis (http://www.ata-atapi.com/. The device driver is compiled using Borland C++ 3.1 and the computer then rebooted under DOS to run the validation tests.

Gordon F. Hughes, Associate Director, U. California San Diego Center for Magnetic Recording Research, gfhughes@ucsd.edu, 858-534-5317, October 2004