

## Re: \*nix data wipe tools

**Source:** <http://www.derkeiler.com/Mailing-Lists/securityfocus/bugtraq/2004-12/0232.html>

---

*Casper.Dik\_at\_Sun.COM*

**Date:** 12/17/04

To: wietse@porcupine.org (Wietse Venema)

Date: Fri, 17 Dec 2004 20:54:13 +0100

>David Cannings:

>> Thomas C. Greene wrote:

>> > I've posted the final versions of a few simple, free shell scripts that i've  
>> > been working on to make data hygiene more convenient on \*nix systems. Thanks  
>> > to list members who helped test them and contributed improvements.

>>

>> Is there any specific advantage of these scripts over bcwipe?

>>

>> [http://www.jetico.com/index.htm#bcwipe\\_unix.htm](http://www.jetico.com/index.htm#bcwipe_unix.htm)

>

>There's a general problem with applications that go through the  
>file system to destroy the contents of a file. Unless one uses  
>very simplistic disk hardware and file systems, there is no guarantee  
>that overwrite requests will actually overwrite the intended bits.  
>For a example, Solaris 10 ZFS uses copy-on-write, to avoid corruption  
>when the system crashes in the middle of an update; many disk drives  
>have write caches built-in so only the last overwrite request takes  
>effect; and non-volatile memory "disks" have a limited number of  
>write cycles and try to avoid hot spots.

I seem to remember that the first version of "PGP wipefile" which was written for DOS also made the assumption that file writes were immediate. No fsync/sync or what not appeared in the Unix port which made it pretty much without effect. (Overwrite it 10 times in the case, then unlink; don't think the kernel will bother to flush the data even once).

ZFS is just one of several "special circumstances" that you may encounter; flash memory devices, e.g., only pretend to be an array of bytes; in actual fact overwriting a single block multiple times will likely cause several different blocks to be overwritten; the original block may not even be overwritten once as the flash memory controller tries to extend the life of the memory by spreading writes evenly around.

We are in fact looking at ways to do "secure delete" as integral part of the ZFS filesystems for those that require it (it obviously will

## SecurityFocus Bugtraq: Re: \*nix data wipe tools

come at a performance price as you will need to do at least 5 I/O ops per block to have some chance of actually erasing the data so it's not likely going to be an "always on feature")

Personally I prefer a "lose key, lose data" approach to secure erase; it's much easier to lose or securely erase a key than it is to do so with a lot of data.

Casper