

Re: hard drive scrubbing utility

Source: <http://www.tech-archive.net/Archive/VC/microsoft.public.vc.mfc/2007-01/msg00686.html>

- *From:* Joseph M. Newcomer <newcomer@xxxxxxxxxxxxx>
 - *Date:* Fri, 12 Jan 2007 22:56:49 -0500
-

Well, because it would have to write the **entire** disk. Have you actually tried to format completely a, say, 300GB drive? It takes a **long** time. So remote-boot systems just initialize the master file block, build the initial directories from a server, and all the contents are virtualized in that the bits themselves aren't loaded onto the machine, but a low-level file driver redirects file requests to the server on-demand. The more sophisticated ones then treat the local file system as a cache, so subsequent operations will use the local copy. But on a reboot, the local copies are discarded and the cache is effectively flushed. Note that none of this actually involves doing more than the **minimum** number of writes to the hard drive. The goal here is to have a boot cycle that happens in under a half-hour.

I have not formatted a hard drive in a long time, and I realize that the speed of hard drives has increased, but not proportional to their density. I used to spend 30 minute formatting a 70MB MFM drive. The smallest modern drive is 1000 times larger, but not 1000 times faster. But a scrubber has to actually **write** every sector of the hard drive. Every sector. Now, you can imagine optimizations where the scrubber writes, in a hidden location on the hard drive, a "high water mark", but there are ways around that. And it would still have to rewrite every sector in the cache, which is often a few gigabytes.

Performance becomes the major issue here. All things are possible if you don't have a time budget. Ultimately, performance matters, sometimes it is the **only** parameter that matters,, and the key is figuring out architectures that make high performance possible. Remote-boot virtualized systems have done this, and the compromise is that the old bits on the disk remain as old bits on the disk; only relevant bits are overwritten. Remote file caching is a very sophisticated technology; for example, you can leave the files in place but do checksum compares against the master copy, so all you need to download from the server is the checksum. This minimizes network bandwidth. If someone tampers with the local copy (e.g., installs a virus), then the checksums don't match and the master copy is reloaded. But at no point are the existing bits on the physical hard drive guaranteed to be overwritten; just the logical structure is redefined. So existing physical bits remain. This is **not** the same as a security-scrub, which attempts to destroy the physical bits.

Some friends of mine worked on the Andrew File System (afs) at CMU, which became IBM's Distributed File System product, and later worked on file systems that allow for disconnected laptops. These all have fascinating problems, and reading these papers is very informative.

Try the following google searches

Re: hard drive scrubbing utility

satyanarayana file system
ousterhout file system

A lot of papers come up. I've seen some of the presentations of these results, and they are quite fascinating. These will give some insight into issues of remote-boot and virtualized file systems; many modern products are based on the ideas of these papers, whether actually derived works or independently developed. Note that absolutely none of these have, as a basic premise, any idea about actually scrubbing disks for deletion of secure information; they deal with issues such as how to synchronize file systems, treat the local machine as a cache for a master file copy, etc.

A disk scrubber must guarantee that every sector that could potentially hold information that was formerly part of a file is overwritten. This means that every sector in free space must be scrubbed as it is freed (and you have to worry about what happens if, in the middle of a scrub-recently-freed blocks, the system crashes, so you have to keep track of unscrubbed free blocks), and you have to make sure that, at the point where you write new data in the file, that any blocks that become implicitly freed are also scrubbed. Then, in the case where you say "kill the contents of this file", you have to make sure that the actual physical sectors that represent the file data are scrubbed. All of this is very complex, and virtually none of it can be done from user space. Since there are several levels of abstraction between the logical file system and the physical file system, you have to make sure you are managing all these interactions correctly. Note that none of these issues arise in a virtualized file system. It doesn't care what physical bits are on the system, as long as the *logical* bits are correct as seen by the consumer of those bits. These are almost completely unrelated problem domains.

joe

On 12 Jan 2007 18:33:57 -0800, "avatar70" <jspalding@xxxxxxxxxx> wrote:

Joe,

Thanks again for the dialogue.

I think we are making some progress in our discussion. Since a virtualized disk boot system is able to write data to a hard drive at the time of system start why couldn't the system also perform the method that I outlined? It seems possible to me at least that it likely could. If it could why then wouldn't the method I outlined at the beginning of this topic possibly work?

I do agree though that a virtualized disk boot system is not what I had intended.

Jason

avatar70 wrote:

I would like to start an open discussion:

Re: hard drive scrubbing utility

I have some ideas as to a method to scrub a hard drive on the fly. Of course my ideas are just that, ideas. My ideas are subject to change and hopefully we are able to come up with a reasonable method reduce the need for the sledge hammer method of data theft.

First, my understanding is that the logical block size is 2048 bytes

Second, although a data can be retrieve for ever using a Transmission Electron Microscope (TEM) or Scanning Tunneling Microscope (STM) ... This is what returned to after selling my IT firm...

Third, who cares! Most people look at porn, or are attorneys, medical offices (Have you ever heard of HIPPA) If one of the medical offices donates a computer and it ever opened a medical file the doctor is screwed. CAN I GET AN AMEN BROTHER?

Fourth, The idea

If a file is written that is smaller than 2048 bytes then the file will fit on one block of a hard drive.

Next write a program that when a user logs off the computers and the computer then reboots its self and when rebooting a program is activated. The program the writes the said file until the hard drive has reached its capacity. When the hard drive reaches its maximum capacity it then sweeps itself 7 times.

What are your thoughts as to this concept? If you use this idea in the future please reference me (Jason Spalding)

Joseph M. Newcomer [MVP]

email: newcomer@xxxxxxxxxxxx

Web: <http://www.flounder.com>

MVP Tips: http://www.flounder.com/mvp_tips.htm

.