Re: Sync option destroys flash!

Source: http://linux.derkeiler.com/Mailing−Lists/Kernel/2005−05/2645.html

---

From: Lennart Sorensen (*lsorense_at_csclub.uwaterloo.ca*)
Date: 05/13/05

```
Date:   Fri, 13 May 2005 13:17:58 −0400
To: "Michael H. Warfield" <mhw@wittsend.com>
```

On Fri, May 13, 2005 at 12:20:06PM −0400, Michael H. Warfield wrote:
> I found this out the hard way. (Kissed one brand new $70 USD 1GB flash
> drive good−bye.) According to the man pages for mount, FAT and VFAT
> file systems ignore the "sync" option. It lies. Maybe it use to be
> true, but it certainly lies now. A simple test can verify this. Mount
> a flash drive with a FAT/VFAT file system without the sync option and
> writes to the drive go very fast. Typing "sync" or unmounting the drive
> afterwards, takes time as the buffered data is written to the flash
> drive. This is as it should be. Mount it with the sync option and
> writes are really REALLY slow (worse than they should be just from
> copying the data through USB) but sync and umount come back immediately
> and result in no additional writing to the drive. [Do the preceding
> with only a few files and less than a few meg of data if you value that
> flash.] So... FAT and VFAT are honoring the sync option. This is very
> VERY bad. It's bad for floppies, it's bad for hard drives, it's FATAL
> for flash drives.

Certainly causes lots of unnecesary writes which flash doesn't like.
Given sync doesn't appear to be the default, whyever would you add sync
to vfat?

> Flash drives have a limited number of write cycles. Many many
> thousands of write cycles, but limited, none the less. They are also
> written in blocks which are much larger than the "sector" size report
> (several K in a physical nand flash block, IRC).
>
> What happens, with the sync option on a VFAT file system, is that the
> FAT tables are getting pounded and over−written over and over and over
> again as each and every block/cluster is allocated while a new file is
> written out. This constant overwriting eventually wears out the first
> block or two of the flash drive.

All the flash I have used do automatic wear leveling. Maybe I have only
used high quality flash media (given I am doing work with embedded
industrial grade gear, that is quite plausible). Of course wear
leveling doesn't mean you aren't doing way more writes than necesary,
but it helps spread the load away from the FAT, which would otherwise

quickly die on most flash cards no matter what system you use for
writing to it.

> I had lost a couple of flash keys previously, without realizing what
> was going on, but what send me off investigating this was when I copied
> a 700 Meg file to a brand new 1G USB 2.0 flash memory key in a USB 2.0
> slot. It took over a half an hour to copy to the drive, which really
> had me wondering WTF was wrong! Then, when I went to use the key, I
> found the first couple of blocks were totally destroyed. Read errors
> immediately upon insertion. Then I started digging and found the
> hotplug / HAL / fstab−sync stuff on Fedora Core was mounting USB drives
> with the "sync" option (if less than 2 Gig). I knew from previous
> experience (CF backup cards in a PDA) that repeated pounding on the FAT
> tables would destroy a flash card with a FAT file system. So I reported
> this on the Fedora list. Someone else noticed that the man pages for
> mount state that FAT and VFAT ignore the sync option. Not so,
> obviously... Copying that 700 Meg file resulted in thousands upon
> thousands upon thousands of writes to the FAT table and backup FAT
> table. It simply blew through all the rewrites for those blocks and
> burned them up. Bye bye flash key...

700M = 1.4M 512byte sectors. I guess if it actually writes one sector
at a time and syncs, and it's not a media with good wear leveling, then
yes that would destroy the sectors holding the FAT. Ouch. Crappy
media, and bad way to treat it. Unfortunately there is no bug, just
user error, and potentially badly designed flash media firmware.

> On a floppy, this would result in an insane amount of jacking around
> back and forth between data sectors and the FAT sectors. In addition to
> taking forever, that would shorten the life of the diskettes and the
> drive itself, but who cares about floppies any more. On a real hard
> drive, this will cause "head resonances" as the heads go through
> constant high speed seeks between the cylinder with the FAT tables and
> the data cylinders. That can't be good, on a continuous basis, for
> drive life. But it's really a disaster for flash memory. It's going to
> cause premature failure in most flash memory, even if it doesn't kill
> them right off as it did in my case with a 700 Meg file.
>
> Can we go back to ignoring "sync" on FAT and VFAT? I can't see where
> it does much good. You might corrupt a file system if you unplugged it
> while dirty but it beats the hell out of physically burning it up and
> destroying the drive!

How about you just don't use the sync option with fat when you don't
mean to use sync? sync does exactly what it should, which just happens
to not be what you want, so don't use it.

> If it's decided that the FAT and VFAT file systems MUST obey the sync
> option then please do something about a special case for the FAT tables!
> Sync the data if thou must buti... Thou shalt not, must not, whack off
> on the FAT tables!!!

Then the sync option wouldn't be much use anymore.

> *Another option would be to only sync the FAT and VFAT file systems upon*
> *close of the file being written or upon close of the last file open on*
> *the file system (fs not busy) but that might not help in the case of a*
> *whole lotta little files...*

Again not very useful then.

> *I'm also going to file a couple of bug reports in bugzilla at RedHat*
> *but this seems to be a more fundamental problem than a RedHat specific*
> *problem. But, IMHO, they should never be setting that damn sync flag*
> *arbitrarily.*

No they certainly should not, but it may have something to do with
making life easier for kde/gnome desktops and automatic mount/umount of
media. Dumb idea still, but that happens sometimes.

Len Sorensen
−