

# Re: mounting LVM partitions fails after etch upgrade

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Debian/2007-05/msg03337.html>

---

- *From:* "David Fuchs" <[php4flicks@xxxxxxxx](mailto:php4flicks@xxxxxxxx)>
  - *Date:* Mon, 21 May 2007 14:26:52 +0200
- 

dear all,

a while back I posted to this list because my file systems on LVM over RAID1 would not mount cleanly anymore after upgrade from sarge to etch. this weekend I had time to poke around in the data on both the disks, and found out what was wrong.

as it turns out, since almost a year, *\*no\** data at all was written to one of the disks!! that didn't stop mdadm from happily reporting that everything with the array was in perfect order, though. I rebooted the system a few times during this period, and not even when assembling the array it complained about anything.

due to the upgrade of mdadm, it seems that the s/w raid started using both disks again, and by writing data to the 'old' disk, corrupting some of the out-of-date data there. I'm glad I didn't try to fix this with fsck, it probably would have completely toasted the data on both disks.

how can such a catastrophic failure of a raid array happen, and worse, go completely unnoticed? I don't think it's a config issue, it perfectly mirrored all data before that point. both disks are physically perfect, not a single bad block.

cheers,  
– Dave.

On 5/6/07, Douglas Allan Tutty <[dtutty@xxxxxxxxxxxxxxxx](mailto:dtutty@xxxxxxxxxxxxxxxx)> wrote:

On Sun, May 06, 2007 at 03:25:02PM +0200, David Fuchs wrote:  
> I have just upgraded my sarge system to etch, following exactly the upgrade  
> instructions at <http://www.us.debian.org/releases/etch/i386/release-notes/>.  
>  
> now my system does not boot correctly anymore... I'm using RAID1 with two  
> disks, / is on md0 and all other mounts (/home/, /var, /usr etc) are on md1  
> using LVM.  
>  
> the first problem is that during boot, only md0 gets started. I can get

## Re: mounting LVM partitions fails after etch upgrade

- > around this by specifying break=mount on the kernel boot line and manually
- > starting md1, but where need I change what so that md1 gets started at this
- > point as well?
- >
- > after manually starting md1 and continuing to boot, I get errors like
- >
- > Inode 184326 has illegal block(s)
- > /var: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY (i.e. without the -a or
- > -o
- > options)
- >
- > ... same for all other partitions on that volume group
- >
- > fsck died with exit status 4
- > A log is being saved in /var/log/fsck/checkfs if that location is
- > writable.(it is not)
- >
- > at this point I get dropped to a maintenance shell. when I select to
- > continue the boot process:

What happens if instead of forcing a boot you do what it says: run fsck without the -a or -o options?

- >
- > EXT3-fs warning: mounting fs with errors. running e2fsck is recommended
- > EXT3 FS on dm-4, internal journal
- > EXT3-FS: mounted filesystem with ordered data mode.
- > ... same for all mounts (same for dm-3, dm-2, dm-1, dm-0)
- >
- > EXT3-fs error (device dm-1) in ext3\_reserve\_inode\_write: Journal has aborted
- > EXT3-fs error (device dm-1) in ext3\_orphan\_write: Journal has aborted
- > EXT3-fs error (device dm-1) in ext3\_orphan\_del: Journal has aborted
- > EXT3-fs error (device dm-1) in ext3\_truncate\_write: Journal has aborted
- > ext3\_abort called.
- > EXT3-fs error (device dm-1): ext3\_journal\_start\_sb: Detected aborte
- > djournal
- > Remounting filesystem read-only
- >
- > and finally I get tons of these:
- >
- > dm-0: rw-9, want=6447188432, limit=10485760
- > attempt to access beyond end of device
- >
- > the system then stops for a long time (~5 minutes) at "starting syslog
- > service" but eventually the login prompt comes up, and I can log in, see all
- > my data, and even (to my surprise) write to the partitions on md1...
- >
- > ...which probably corrupts the fs even more.
  
- > what the hell is going on here? thanks a lot in advance for any help!
- >

## Re: mounting LVM partitions fails after etch upgrade

What is going on is that you started with a simple booting error that has propagated into filesystem errors. Those errors are compounded by forcing a mount of a filesystem with errors. Remember that the system that starts LVM and raid itself exists on the disks....

What you need is a shell with the root fs either totally unmounted or mounted ro. Does booting single-user work? What about telling the kernel `init=/bin/sh`? From there, you can check the status of the mds with:

```
#!/sbin/mdadm -D /dev/md0
#!/sbin/mdadm -D /dev/md1
...
```

check the status of the logical volumes:  
`#!/sbin/lvdisplay [lvname]`

and then check the filesystems with:

```
#!/sbin/e2fsck -f -c -c /dev/...
```

Only once you get the filesystems fully functional should you attempt to boot further.

Doug.

--

To UNSUBSCRIBE, email to `debian-user-REQUEST@xxxxxxxxxxxxxxxxxx` with a subject of "unsubscribe". Trouble? Contact `listmaster@xxxxxxxxxxxxxxxxxx`

--

To UNSUBSCRIBE, email to `debian-user-REQUEST@xxxxxxxxxxxxxxxxxx` with a subject of "unsubscribe". Trouble? Contact `listmaster@xxxxxxxxxxxxxxxxxx`