

FAQ: How to fix a degraded filesystem?

Contributed by Michael Felt

I ran into this by chance, and am writing it up before they all go away - forever.

A degraded filesystem seems to be one that does not have all the meta data agreeing. To protect the data that is there the filesystem is mounted, but in ro,degraded status (read-only, degraded).

```
michael@x054:[/home/michael]mount | grep export
```

```
/dev/lvExport /export      jfs2  Mar 12 21:16 ro,degraded
/dev/lvLpp    /export/lpp    jfs2  Mar 12 21:16 ro,degraded
/dev/lvSpot   /export/spot   jfs2  Mar 12 21:16 ro,degraded
```

p.s. the problem seems to be an iscsi "delivery" issue - some/all communication with the iscsi (read NAS) server has been lost. See more below.

Now that this has come back again (borrowed an old NAS device :wink:) I can verify that the solution path is fairly simple. Just unmount and mount again. (because the jfs2log, VGDA, etc gets read again, and the data "loss" is resolved). The hardest part may be getting the filesystem to unmount.

```
michael@x054:[/home/michael]unmount -f /export/spot
```

```
michael@x054:[/home/michael]fsck -n /export/spot
```

The current volume is: /dev/lvSpot

Primary superblock is valid.

*** Phase 1 - Initial inode scan

*** Phase 2 - Process remaining directories

*** Phase 3 - Process remaining files

*** Phase 4 - Check inode allocation map

*** Phase 5 - Check block allocation map

File system is clean.

File system is clean but is marked dirty.

Run fsck with -p parameter to fix.

```
michael@x054:[/home/michael]fsck -p /dev/lvSpot
```

The current volume is: /dev/lvSpot

Open volume exclusive read or write returned, rc = 16

fsck: 0507-289 Device unavailable or locked by another process.

Cannot continue.

```
# chfs -A no /export/spot
```

as well as for the other degraded filesystems and forced that volumn group off, then on again. As I became more experienced in this (the things you can learn from a failing iscsi system) the real key seems to be to offline the iscsi connection and connect the disk again. (easy way is `#rmdev -l hdisk# --` do not bother with the `-d` then you can use `mkdev -l hdisk#` to bring it back online. If you use `#rmdev -dl hdisk#` you will need to use `# cfmgr` to get the disk discovered again.)

If this is iscsi, and your disk does not reconnect immediately you can try my "trick", change the name of the iscsi disk/volume to something else and then change it back again. For me this seems to reset whatever is broken in terms of communication. FYI - the information on the disks has ALWAYS been okay.

The switch to degraded,ro seems to come fast enough that no damage can be done. At least I have not lost any data.

The ultimate solution - new NAS firmware, and if that fails, a better NAS device.