

Re: Bad blocks not getting remapped on Maxtor drive

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.misc/2004-03/2631.html>

From: P Gentry (rdgentry1_at_cablelynx.com)

Date: 03/28/04

Date: 28 Mar 2004 09:31:38 -0800

John Hendriks <hjohn@remove.this.xs4all.nl> wrote in message
news:<4066adc1\$0\$569\$e4fe514c@news.xs4all.nl>...

> Hiya all,

>

> I've been trying to fix a problem with bad blocks on a fairly new Maxtor
> 300 GB drive (only about 6 months old and not used that much).

> /sbin/badblocks will report about 100 bad blocks in a single range, and

> keeps reporting them even when I use the `-n` or `-w` options to write to

> those blocks -- I was under the impression a disk will remap bad blocks

> when written to:

>

> `# /sbin/badblocks -w -c 16384 /dev/hdg2 13120000 13110000`

>

> -> Results in about 100 bad blocks to be printed (takes about 3 minutes

> or so... about 1 in 20 times I try this, the drive will make the

> parking sound *clunk*).

>

> I've also tried zeroing that area with `dd`, like this:

>

> `# dd if=/dev/zero of=/dev/hdg bs=512 seek=346379385 count=20000`

> `dd: writing '/dev/hdg': Input/output error`

> `11152+0 records in`

> `11151+0 records out`

>

> I did the above atleast 20 times with both programs, results have been the

> same. The output from `/sbin/badblocks` varies slightly though, sometimes

> printing the blocks in a different order or omitting a few at the start/end

> of the range. Each time it reported around 100 bad blocks though.

>

> I also tried using the `-D` option of `hdparm` to activate defect management,

> however this option only resulted in an error for the Maxtor drive (I'm

> assuming the drive doesn't offer an option to change this and defaults

> to on).

>

> All of this seems to point out that the drive has run out of spare

> blocks which it can use to remap. However, `smartctl` reports it only has

comp.os.linux.misc: Re: Bad blocks not getting remapped on Maxtor drive

- > 2 sectors reallocated. After several hours of writing to that area and
- > trying all sorts of things, I managed to get it to 3 sectors
- > reallocated, but it hasn't changed since then.

Note the differences --- bocks vs sectors --- ie., fs and physical disk layout. SMART is pretty conservative in "auto repair" mode (aka "normal use") and you'll need Maxtor's disk utils to get more aggressive in repair attempts. See below.

- >
- > Here's part of the smartctl -a output:
- [snip]
- > <cut ATA error log>
- >
- > SMART Self-test log structure revision number 1
- > Num Test_Description Status Remaining LifeTime(hours) LBA_of_first_error
- > # 1 Short offline Completed: read failure 60% 4967 0x14a5800c
- > # 2 Extended captive Completed: read failure 20% 4963 0x14a5800c
- >
- > Note 1:
- > Everytime I try to write to the defect area, the numbers Run_Out_Cancel
- > and Hardware_ECC_Recovered go up, as well as the amount of ATA errors in
- > the log. The Current_Pending_Sector value changed from 2 to 3 at the same
- > time when Reallocated_Sector_Ct changed to 3.

Sounding ominous --- SMART has limited capacity for relocating bad sectors to "good, reserved" sectors on disk

- > Note 2:
- > Run_Out_Cancel might be the incorrect name for Maxtor drives.

Makers vary in how they "deviate" from the standard --- another reason to use the maker's disk utils for thorough checking/repair.

- > As far as I can see, the drive is still perfectly fine, it just won't remap
- > the defective blocks. ...

This is "normal" behavior. SMART knows it can't accurately _read_ the sector, so why would it _automatically_ try to read it and re-write it? You need maker's utils to "force" this attempted repair. Also blocks your attempts to manipulate those sectors in any way through the OS.

- > ... I've considered remapping the blocks in software (with
- > mke2fs -c) but since this drive is going to be part of a raid-array I would
- > like to be able to rely on the drive to remap blocks on its own.

This won't help and may result in a _totally_ trashed disk --- DON'T!

- > Any help would be appreciated!
- >

Re: Bad blocks not getting remapped on Maxtor drive

> --John

First ...!!!!

Do not write to this hard disk at all till you've confirmed the condition of the drive. See below.

Second ... !!!!!

If you have any data you don't care to lose, back it up now!

Third ... !!!!!!!

Don't change any physical settings/connections on this drive till you've run the suggested test/recovery routines.

You need to visit the Maxtor web site and get a copy of their PowerMax utils -- their site doesn't like my Konqueror settings and I'm too lazy to pull it up in Mozilla to navigate to the right page for you. IIRC it's not hard to find, but try to get the separate utils if you can, as you will need a DOS sys disk to boot up the PowerMax software (I think this DOS sys disk requirement is still true).

While there pick up a current list of error codes, check KB and FAQ for insights, and review procedures necessary to obtain an RMA if disk proves bad.

Sounds like you are getting early(?) signs of a bad hard drive. The S.M.A.R.T. routines are somewhat limited in what they will attempt automatically in recovering/relocating bad sectors -- ie., sectors that cannot be read correctly. And it sounds like you have a bunch.

Visit the smartmon tools site to get more complete, up-to-date info re: error messages and brand peculiarities:

<http://smartmontools.sourceforge.net/>

and maybe a look at this article by the author:

<http://www.linuxjournal.com/article.php?sid=6983>

I've just finished refreshing my skills in this area -- the WD disk proved to be toast -- and the fact is that smartctl and tools cannot "force" recovery attempts like the manufacturers' disk utilities can. If the drive is going bad, Maxtor will require you to run test/recovery and post results to them before issuing an RMA anyway.

If I were you, I would start polishing my "I demand a new replacement disk" speech -- especially as you are wanting to place this drive in a RAID. The maker marks the bad sectors at time of manufacture and you really should not be getting any new ones this early in the disk's life -- unless it's been subjected to physical/electrical shocks or passed near a degaussing device.

BTW, you did get a new 40 pin/80 wire ide cable with your drive, didn't you? It's definitely required for reliability. CTTOI, you should have received a PowerMax disk also (but check site in any case

comp.os.linux.misc: Re: Bad blocks not getting remapped on Maxtor drive

for newer version, etc.)

hth,

prg

email above disabled