

A little story of failed raid5 (3ware 8000 series)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2007-08/msg00269.html>

- *From:* "Artem Kuchin" <matrix@xxxxxxxxxxxx>
 - *Date:* Tue, 21 Aug 2007 02:38:34 +0400
-

Hello!

Here is the newest story of mine about how one should never use raid5.

Controller is 8xxx-4LP.
I have a simple 360GB raid5 with 4 drives since 2004.
Only about a year ago i realized how much speed i have wasted be saving lousy 120GB. I should have choosen bigger driver and setup two mirrors instead.

But that's no the point. A week ago one driver just totally failed. It fell out of the unit and when i tried to rebuild the unit it failed. It seemed like the driver electronis failed. ANyhow, i have found newest 160gb seagate driver for replacement (twice as thin, very nicely done electornics on it).

A day ago at 11 am i have turn off the server, pull out the old driver, installed a new one, turned of the server and started rebuild in an hour from remote location via web interface. After about 5 minuted the machine became unresponsive. Tried rebooting – nothing. I went to the machine and fingure out, that rebuild failed (0%) and some data cannot be read because of bad sectors.

Well, hell, i thoght. Maybe i could tell teh controller to ignore all the errors and just some rebuilding and the figure out which driver failed, replace it, rebuild again and restore corrupted data from backup. Noway, controller said.

- i cannot make it ignore read errors
- i cannot figure out which driver has bad sectors (maybe someone know it?)

But i don't understand how and why it happened. ONLY 6 hours ago (a night before) all those files were backed up fine w/o any read error. And now, right after replacing the driver and starting rebuild it said that there are bad sectors all over those file. How come?

A little story of failed raid5 (3ware 8000 series)

Well. Since we have a buch of full and inceremnetal paraoid backups no data was lost and we are in process of recovering. However, i simply imaged what would happed if one more driver completelly failed. That would mean that we have lost all data, since any of the disk which left do not contain any readable copy of one data (unlink mirror, for example).

So, we are migrating to mirror config with huge disks.

I am thinking about raid10 for more perfomance. It seems a lot more safe, since if any pair of disks failed the data is still readable and even if all disks have bad block the data can be easily recovered by fairly simply script from the couterpart. But the problem, however,

So, no raid5 or even raid 6 for me any more. Never!

--

Regards,
Artem

freebsd-stable@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-stable>

To unsubscribe, send any mail to "freebsd-stable-unsubscribe@xxxxxxxxxxx"