

Re: "ad0: TIMEOUT – WRITE_DMA" type errors with 7.0–RC1

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2008-01/msg00633.html>

- *From:* Joe Peterson <joe@xxxxxxxxxxxxxxxxxxx>
 - *Date:* Fri, 25 Jan 2008 12:15:58 –0700
-

Jeremy Chadwick wrote:

What you've shown is usually the sign of a disk-related problem. It's very obvious when it's just one disk reporting DMA errors. You use ZFS, so chances are you have more than one disk in a pool/volume -- there's no indication ad1, ad4, ad6, etc. are failing, so this seems to indicate something specific to ad0.

Jeremy, thanks for the response – I have tried to answer all of your questions below...

In my case, I am using only one disk (ad0) for FreeBSD, and I am only using one partition on this disk in my ZFS pool. So, in this case, unfortunately, it's not possible to tell from the fact that only ad0 is listed that it is specific to this drive.

Manufacturers pick very passive (non-aggressive) thresholds for error conditions on disks, so disks which are failing very commonly show "PASSED" during SMART analysis. To make matters worse, most users I know read SMART stats incorrectly (they're easy to misinterpret).

Yep, I am also always skeptical of smart reports. That's one reason I am very interested in ZFS. I don't trust the drive to be completely reliable, and the fact that ZFS does end-to-end data integrity is very intriguing.

Can you please provide output of the following:

```
* smartctl -a /dev/ad0
```

OK, I've attached this to the end of this email.

Re: "ad0: TIMEOUT – WRITE_DMA" type errors with 7.0–RC1

* atacontrol cap ad0

Protocol ATA/ATAPI revision 7
device model ST3500630A
serial number 9QG0DG03
firmware revision 3.AAE
cylinders 16383
heads 16
sectors/track 63
lba supported 268435455 sectors
lba48 supported 976773168 sectors
dma supported
overlap not supported

Feature Support Enable Value Vendor
write cache yes yes
read ahead yes yes
Tagged Command Queuing (TCQ) no no 0/0x00
SMART yes yes
microcode download yes yes
security yes no
power management yes yes
advanced power management no no 65278/0xFEFE
automatic acoustic management no no 0/0x00 208/0xD0

* atacontrol info <ata0, ata1, etc. — any controller used by ZFS>

Master: ad0 <ST3500630A/3.AAE> ATA/ATAPI revision 7
Slave: ad1 <ST3160812A/3.AAH> ATA/ATAPI revision 7

(but note that ad1 is not used by FreeBSD)

* Relevant dmesg output that indicates what kind of ATA controller these disks are attached to. Start with output from 'ad0:' and work backwards. For example, ad0 on this machine is using an Intel ICH6 controller:
atapci0: <Intel ICH6 SATA150 controller> port
0x1f0–0x1f7,0x3f6,0x170–0x177,0x376,0xf000–0xf00f at device 31.2 on pci0
ata0: <ATA channel 0> on atapci0
ad0: 238475MB <WDC WD2500KS–00MJB0 02.01C03> at ata0–master SATA150

atapci0: <Intel ICH4 UDMA100 controller> port
0x1f0–0x1f7,0x3f6,0x170–0x177,0x376,0xf000–0xf00f at device 31.1 on pci0
ata0: <ATA channel 0> on atapci0

Re: "ad0: TIMEOUT – WRITE_DMA" type errors with 7.0–RC1

Re: "ad0: TIMEOUT – WRITE_DMA" type errors with 7.0–RC1

ata0: [ITHREAD]

ad0: 476940MB <Seagate ST3500630A 3.AAE> at ata0–master UDMA100

SMART stats which are labelled "Offline" are only updated when a short or long offline test is performed. Have you tried using "smartctl –t short /dev/ad0" and "smartctl –t long /dev/ad0" to see if any of the raw values on the far right column increment?

I just tried one:

1 Short offline Completed without error 00% 5252

–

2 Short offline Completed without error 00% 5252

–

Also, none of the numbers that were zero incremented, esp:

198 Offline_Uncorrectable 0x0010 100 100 000 Old_age

Offline – 0

Also, no more errors were reported in the system log during the self–tests.

Have you tried using "zpool scrub" on the ZFS pool, then "zpool status" to see if READ/WRITE/CHKSUM counters increment or if the "scrub" line states there were errors?

OK, I started a scrub, and it will take some more time to complete... But I get the following with status. Could this be due to the timeouts and failures? I suspect so, so maybe this is not surprizing. I'd also guess that this doesn't necessarily point to the drive, but anything in the chain of events... I do not have a mirror or RADI–Z, so I guess the reason there was "no data loss" (yet) is because the checksum passed, and maybe it just had to retry...? Anyway, here's the output so far:

pool: tank

state: ONLINE

status: One or more devices has experienced an unrecoverable error. An attempt was made to correct the error. Applications are unaffected.

action: Determine if the device needs to be replaced, and clear the errors using 'zpool clear' or replace the device with 'zpool replace'.

see: <http://www.sun.com/msg/ZFS–8000–9P>

scrub: scrub in progress, 2.50% done, 1h58m to go

config:

NAME STATE READ WRITE CKSUM

tank ONLINE 1 3 0

ad0s1d ONLINE 1 3 0

Re: "ad0: TIMEOUT – WRITE_DMA" type errors with 7.0–RC1

Re: "ad0: TIMEOUT – WRITE_DMA" type errors with 7.0–RC1

errors: No known data errors

Other things which have fixed problems in the past for others:

- * BIOS updates
- * Change of motherboards (sometimes replacing board with same model, other times going with a completely different vendor (implies weird implementation issues or BIOS problems))

I've been using this same motherboard/BIOS for a long time (as well as this drive), so no changes have happened to the HW recently. The BIOS is the newest, available, I believe (It's a Tyan Trinity S2099, so it's a few years old)

- * Changing SATA cables

I'm using regular ATA 80–pin cables. Also, these seem to have been working fine for quite a while now. But, yes, I have also witnessed bad cable issues on older systems in the past. I certainly could try a new cable and see if it helps.

- * Getting a larger power supply (usually when lots of disk are involved)

I only have two drives, so I think the PS has enough capacity in my case.

Anyway, thanks for the reply and further questions. Let me know if anything I've sent back is helpful!

Thanks, Joe

smartctl version 5.37 [i386–portbld–freebsd7.0] Copyright (C) 2002–6 Bruce Allen
Home page is <http://smartmontools.sourceforge.net/>

=== START OF INFORMATION SECTION ===

Model Family: Seagate Barracuda 7200.10 family

Device Model: ST3500630A

Serial Number: 9QG0DG03

Firmware Version: 3.AAE

User Capacity: 500,107,862,016 bytes

Device is: In smartctl database [for details use: –P show]

ATA Version is: 7

ATA Standard is: Exact ATA specification draft version not indicated

Local Time is: Fri Jan 25 09:55:13 2008 MST

SMART support is: Available – device has SMART capability.

SMART support is: Enabled

Re: "ad0: TIMEOUT – WRITE_DMA" type errors with 7.0–RC1

=== START OF READ SMART DATA SECTION ===
SMART overall–health self–assessment test result: PASSED

General SMART Values:

Offline data collection status: (0x82) Offline data collection activity was completed without error.
Auto Offline Data Collection: Enabled.
Self–test execution status: (0) The previous self–test routine completed without error or no self–test has ever been run.
Total time to complete Offline data collection: (430) seconds.
Offline data collection capabilities: (0x5b) SMART execute Offline immediate.
Auto Offline data collection on/off support.
Suspend Offline collection upon new command.
Offline surface scan supported.
Self–test supported.
No Conveyance Self–test supported.
Selective Self–test supported.
SMART capabilities: (0x0003) Saves SMART data before entering power–saving mode.
Supports SMART auto save timer.
Error logging capability: (0x01) Error logging supported.
General Purpose Logging supported.
Short self–test routine recommended polling time: (1) minutes.
Extended self–test routine recommended polling time: (163) minutes.

SMART Attributes Data Structure revision number: 10

Vendor Specific SMART Attributes with Thresholds:

ID#	ATTRIBUTE_NAME	FLAG	VALUE	WORST	THRESH	TYPE	UPDATED	WHEN_FAILED	RAW_VALUE
1	Raw_Read_Error_Rate	0x000f	114 071 006	Pre–fail	Always	–	82422948		
3	Spin_Up_Time	0x0003	093 093 000	Pre–fail	Always	–	0		
4	Start_Stop_Count	0x0032	100 100 020	Old_age	Always	–	56		
5	Reallocated_Sector_Ct	0x0033	100 100 036	Pre–fail	Always	–	1		
7	Seek_Error_Rate	0x000f	084 060 030	Pre–fail	Always	–	286126605		
9	Power_On_Hours	0x0032	095 095 000	Old_age	Always	–	5250		
10	Spin_Retry_Count	0x0013	100 100 097	Pre–fail	Always	–	0		
12	Power_Cycle_Count	0x0032	100 100 020	Old_age	Always	–	59		
187	Unknown_Attribute	0x0032	100 100 000	Old_age	Always	–	0		
189	Unknown_Attribute	0x003a	100 100 000	Old_age	Always	–	0		
190	Temperature_Celsius	0x0022	065 056 045	Old_age	Always	–	605749283		
194	Temperature_Celsius	0x0022	035 044 000	Old_age	Always	–	35 (Lifetime Min/Max 0/15)		
195	Hardware_ECC_Recovered	0x001a	063 046 000	Old_age	Always	–	166181300		
197	Current_Pending_Sector	0x0012	100 100 000	Old_age	Always	–	0		
198	Offline_Uncorrectable	0x0010	100 100 000	Old_age	Offline	–	0		

Re: "ad0: TIMEOUT – WRITE_DMA" type errors with 7.0-RC1

199 UDMA_CRC_Error_Count 0x003e 200 200 000 Old_age Always – 0
200 Multi_Zone_Error_Rate 0x0000 100 253 000 Old_age Offline – 0
202 TA_Increase_Count 0x0032 100 253 000 Old_age Always – 0

SMART Error Log Version: 1
No Errors Logged

SMART Self-test log structure revision number 1

SMART Selective self-test log data structure revision number 1

SPAN MIN_LBA MAX_LBA CURRENT_TEST_STATUS

1 0 0 Not_testing

2 0 0 Not_testing

3 0 0 Not_testing

4 0 0 Not_testing

5 0 0 Not_testing

Selective self-test flags (0x0):

After scanning selected spans, do NOT read-scan remainder of disk.

If Selective self-test is pending on power-up, resume after 0 minute delay.

freebsd-stable@xxxxxxxxxxx mailing list

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

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