
Subject: SATA HDD Problem

Posted by [Markus Hardiyanto](#) on Mon, 16 Jul 2007 08:36:27 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello,

i installed openvz with 2.6.18 kernel and having problem with sata HDD on my server.. here is the error message from /var/log/messages:

Jul 14 06:46:39 cl-44 kernel: ata1.00: exception Emask 0x0 SAct 0x0 SErr 0x0 action 0x0

Jul 14 06:46:39 cl-44 kernel: ata1.00: tag 0 cmd 0xb0 Emask 0x1 stat 0x51 err 0x4 (device error)

Jul 14 06:46:39 cl-44 kernel: ata1: EH complete

Jul 14 06:46:39 cl-44 kernel: ata1.00: exception Emask 0x0 SAct 0x0 SErr 0x0 action 0x0

Jul 14 06:46:39 cl-44 kernel: ata1.00: tag 0 cmd 0xb0 Emask 0x1 stat 0x51 err 0x4 (device error)

Jul 14 06:46:39 cl-44 kernel: ata1: EH complete

Jul 14 06:46:39 cl-44 kernel: ata1.00: exception Emask 0x0 SAct 0x0 SErr 0x0 action 0x0

Jul 14 06:46:39 cl-44 kernel: ata1.00: tag 0 cmd 0xb0 Emask 0x1 stat 0x51 err 0x4 (device error)

Jul 14 06:46:39 cl-44 kernel: ata1: EH complete

Jul 14 06:46:39 cl-44 kernel: ata1.00: exception Emask 0x0 SAct 0x0 SErr 0x0 action 0x0

Jul 14 06:46:39 cl-44 kernel: ata1.00: tag 0 cmd 0xb0 Emask 0x1 stat 0x51 err 0x4 (device error)

Jul 14 06:46:39 cl-44 kernel: ata1: EH complete

Jul 14 06:46:39 cl-44 kernel: ata1.00: exception Emask 0x0 SAct 0x0 SErr 0x0 action 0x0

Jul 14 06:46:39 cl-44 kernel: ata1.00: tag 0 cmd 0xb0 Emask 0x1 stat 0x51 err 0x4 (device error)

Jul 14 06:46:39 cl-44 kernel: ata1: EH complete

Jul 14 06:46:39 cl-44 kernel: ata1.00: exception Emask 0x0 SAct 0x0 SErr 0x0 action 0x0

Jul 14 06:46:39 cl-44 kernel: ata1.00: tag 0 cmd 0xb0 Emask 0x1 stat 0x51 err 0x4 (device error)

Jul 14 06:46:39 cl-44 kernel: ata1: EH complete

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: ATAPI reset complete

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: ATAPI reset complete

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: ATAPI reset complete

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: ATAPI reset complete

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: ATAPI reset complete

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: ATAPI reset complete

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:50:13 cl-44 kernel: hda: status error: status=0x20 { DeviceFault }

Jul 14 06:50:13 cl-44 kernel: ide: failed opcode was: unknown

Jul 14 06:55:38 cl-44 smartd[2673]: smartd version 5.36
[x86_64-redhat-linux-gnu] Copyright (C) 2002-6 Bruce Allen

Jul 14 06:55:38 cl-44 smartd[2673]: Home page is
<http://smartmontools.sourceforge.net/>

Jul 14 06:55:38 cl-44 smartd[2673]: Opened configuration file
/etc/smartd.conf

Jul 14 06:55:38 cl-44 smartd[2673]: Configuration file /etc/smartd.conf
parsed.

Jul 14 06:55:38 cl-44 smartd[2673]: Device: /dev/sda, opened

Jul 14 06:55:38 cl-44 smartd[2673]: Device: /dev/sda, not found in
smartd database.

Jul 14 06:55:39 cl-44 smartd[2673]: Device: /dev/sda, is SMART capable.
Adding to "monitor" list.

Jul 14 06:55:39 cl-44 smartd[2673]: Monitoring 1 ATA and 0 SCSI devices

Jul 14 06:55:39 cl-44 smartd[2673]: Device: /dev/sda, 60 Currently
unreadable (pending) sectors

Jul 14 06:55:39 cl-44 smartd[2673]: Sending warning via mail to root ...

Jul 14 06:55:39 cl-44 smartd[2673]: Warning via mail to root: successful

Jul 14 06:55:39 cl-44 smartd[2673]: Device: /dev/sda, 65 Offline
uncorrectable sectors

Jul 14 06:55:39 cl-44 smartd[2673]: Sending warning via mail to root ...

Jul 14 06:55:39 cl-44 smartd[2673]: Warning via mail to root: successful

Jul 14 06:55:39 cl-44 smartd[2687]: smartd has fork()ed into background
mode. New PID=2687.

Jul 14 07:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently
unreadable (pending) sectors

Jul 14 07:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline
uncorrectable sectors

Jul 14 07:55:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently
unreadable (pending) sectors

Jul 14 07:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline
uncorrectable sectors

Jul 14 08:04:49 cl-44 init: Trying to re-exec init

Jul 14 08:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently
unreadable (pending) sectors

Jul 14 08:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline
uncorrectable sectors

Jul 14 08:55:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently
unreadable (pending) sectors

Jul 14 08:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline
uncorrectable sectors

Jul 14 09:25:40 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently
unreadable (pending) sectors

Jul 14 09:25:40 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline
uncorrectable sectors

Jul 14 09:55:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently
unreadable (pending) sectors

Jul 14 09:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline
uncorrectable sectors

Jul 14 10:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 10:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 10:55:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 10:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 11:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 11:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 11:55:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 11:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 12:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 12:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 12:55:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 12:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 13:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 13:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 13:55:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 13:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 14:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 14:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 14:55:40 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 14:55:40 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 15:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 15:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 15:55:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 15:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 16:25:39 cl-44 smartd[2687]: Device: /dev/sda, 60 Currently unreadable (pending) sectors

Jul 14 16:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 16:55:39 cl-44 smartd[2687]: Device: /dev/sda, 63 Currently unreadable (pending) sectors

Jul 14 16:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 17:25:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Currently unreadable (pending) sectors

Jul 14 17:25:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 17:55:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Currently unreadable (pending) sectors

Jul 14 17:55:39 cl-44 smartd[2687]: Device: /dev/sda, 65 Offline uncorrectable sectors

Jul 14 18:25:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Currently unreadable (pending) sectors

Jul 14 18:25:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Offline uncorrectable sectors

Jul 14 18:55:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Currently unreadable (pending) sectors

Jul 14 18:55:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Offline uncorrectable sectors

Jul 14 19:25:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Currently unreadable (pending) sectors

Jul 14 19:25:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Offline uncorrectable sectors

Jul 14 19:55:40 cl-44 smartd[2687]: Device: /dev/sda, 90 Currently unreadable (pending) sectors

Jul 14 19:55:40 cl-44 smartd[2687]: Device: /dev/sda, 90 Offline uncorrectable sectors

Jul 14 20:25:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Currently unreadable (pending) sectors

Jul 14 20:25:39 cl-44 smartd[2687]: Device: /dev/sda, 90 Offline uncorrectable sectors

from smartctl command:

```
# smartctl -l error -d ata /dev/sda
```

smartctl version 5.36 [i686-redhat-linux-gnu] Copyright (C) 2002-6 Bruce Allen

Home page is <http://smartmontools.sourceforge.net/>

=== START OF READ SMART DATA SECTION ===

SMART Error Log Version: 1

ATA Error Count: 52 (device log contains only the most recent five errors)

CR = Command Register [HEX]

FR = Features Register [HEX]

SC = Sector Count Register [HEX]

SN = Sector Number Register [HEX]

CL = Cylinder Low Register [HEX]

CH = Cylinder High Register [HEX]

DH = Device/Head Register [HEX]

DC = Device Command Register [HEX]

ER = Error register [HEX]

ST = Status register [HEX]

Powered_Up_Time is measured from power on, and printed as

DDd+hh:mm:SS.sss where DD=days, hh=hours, mm=minutes,

SS=sec, and sss=millisec. It "wraps" after 49.710 days.

Error 52 occurred at disk power-on lifetime: 4015 hours (167 days + 7 hours)

When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:

ER ST SC SN CL CH DH

-- -- -- -- --

10 51 01 6e 43 f9 ed

Commands leading to the command that caused the error were:

CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name

```
37 00 01 6e 43 f9 ed 00    05:03:57.209    SET MAX ADDRESS EXT
27 00 00 6e 43 f9 e0 00    05:03:57.164    READ NATIVE MAX ADDRESS EXT
37 00 00 6e 43 f9 ed 00    05:03:57.164    SET MAX ADDRESS EXT
27 00 01 6e 43 f9 e0 00    05:03:57.164    READ NATIVE MAX ADDRESS EXT
37 00 01 6e 43 f9 ed 00    05:04:01.507    SET MAX ADDRESS EXT
```

Error 51 occurred at disk power-on lifetime: 4015 hours (167 days + 7 hours)

When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:

ER ST SC SN CL CH DH

10 51 01 6e 43 f9 ed

Commands leading to the command that caused the error were:

CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name

```
37 00 01 6e 43 f9 ed 00    05:03:57.209    SET MAX ADDRESS EXT
27 00 00 6e 43 f9 e0 00    05:03:57.164    READ NATIVE MAX ADDRESS EXT
37 00 00 6e 43 f9 ed 00    05:03:57.164    SET MAX ADDRESS EXT
27 00 01 af 4b f9 e0 00    05:03:57.164    READ NATIVE MAX ADDRESS EXT
```

29 00 01 af 4b f9 e0 00 05:03:57.164 READ MULTIPLE EXT

Error 50 occurred at disk power-on lifetime: 4014 hours (167 days + 6 hours)

When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:

ER ST SC SN CL CH DH

-- -- -- -- --

10 51 01 6e 43 f9 ed

Commands leading to the command that caused the error were:

CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name

-- -- -- -- -- ----- -----

37 00 01 6e 43 f9 ed 00	03:54:04.428	SET MAX ADDRESS EXT
27 00 00 6e 43 f9 e0 00	03:54:02.348	READ NATIVE MAX ADDRESS EXT
37 00 00 6e 43 f9 ed 00	03:54:02.288	SET MAX ADDRESS EXT
27 00 01 6e 43 f9 e0 00	03:54:02.072	READ NATIVE MAX ADDRESS EXT
37 00 01 6e 43 f9 ed 00	03:54:02.001	SET MAX ADDRESS EXT

Error 49 occurred at disk power-on lifetime: 4014 hours (167 days + 6 hours)

When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:

ER ST SC SN CL CH DH

10 51 01 6e 43 f9 ed

Commands leading to the command that caused the error were:

CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name

37 00 01 6e 43 f9 ed 00	03:54:00.392	SET MAX ADDRESS EXT
27 00 00 6e 43 f9 e0 00	03:54:02.348	READ NATIVE MAX ADDRESS EXT
37 00 00 6e 43 f9 ed 00	03:54:02.288	SET MAX ADDRESS EXT
27 00 01 af 4b f9 e0 00	03:54:02.072	READ NATIVE MAX ADDRESS EXT
29 00 01 af 4b f9 e0 00	03:54:02.001	READ MULTIPLE EXT

Error 48 occurred at disk power-on lifetime: 4013 hours (167 days + 5 hours)

When the command that caused the error occurred, the device was active or idle.

After command completion occurred, registers were:

ER ST SC SN CL CH DH

10 51 01 6e 43 f9 ed

Commands leading to the command that caused the error were:

CR FR SC SN CL CH DH DC Powered_Up_Time Command/Feature_Name

37 00 01 6e 43 f9 ed 00 03:24:05.625 SET MAX ADDRESS EXT
27 00 00 6e 43 f9 e0 00 03:24:05.565 READ NATIVE MAX ADDRESS EXT
37 00 00 6e 43 f9 ed 00 03:24:05.349 SET MAX ADDRESS EXT
27 00 01 6e 43 f9 e0 00 03:24:05.278 READ NATIVE MAX ADDRESS EXT
37 00 01 6e 43 f9 ed 00 03:24:05.132 SET MAX ADDRESS EXT

i check on this.. http://bugzilla.kernel.org/show_bug.cgi?id=8650 it seems that it the same problem that i encounter.. how to solve this?

thanks
Best Regards,
Markus

Send instant messages to your online friends <http://uk.messenger.yahoo.com>

Subject: Re: SATA HDD Problem
Posted by [Gregor Mosheh](#) on Mon, 16 Jul 2007 08:55:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Markus Hardiyanto wrote:

> i check on this.. http://bugzilla.kernel.org/show_bug.cgi?id=8650 it seems that it the same problem that i encounter.. how to solve this?

Ouch, that's a rough one!

The bugzilla shows that they're actively working on it as of yesterday, and it sounds as if they're some ways away from a solution. Once they figure it out, it may be some time before the fix makes it into the official Linux kernel, and longer still until it gets into a OpenVZ kernel.

If you're okay with applying the OpenVZ patches to your own kernel source and building your own OpenVZ-capable kernel, your best bet may be to sit tight and keep watching that bugzilla page until they fix this bug in the kernel. Then grab that newest source and give it a try...

Sorry I don't have anything more constructive to say, but that seems to be the status at the moment. :(

--

Gregor Mosheh / Greg Allensworth
System Administrator, HostGIS cartographic development & hosting services
<http://www.HostGIS.com/>

"Remember that no one cares if you can back up,
only if you can restore." - AMANDA

Subject: Re: SATA HDD Problem
Posted by [Markus Hardiyanto](#) on Mon, 16 Jul 2007 09:23:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

yeah.. its kind of big problem.. i already bought the hardware and seems can't replace it.
but some source said that the problem doesn't occur on centos 4.5 which has 2.6.9 kernel, has
anyone tried this?

thanks

Best Regards,
Markus

----- Original Message -----

From: Gregor Mosheh <gregor@hostgis.com>
To: users@openvz.org
Sent: Monday, July 16, 2007 3:55:29 PM
Subject: Re: [Users] SATA HDD Problem

Markus Hardiyanto wrote:

> i check on this.. http://bugzilla.kernel.org/show_bug.cgi?id=8650 it seems that it the same
problem that i encounter.. how to solve this?

Ouch, that's a rough one!

The bugzilla shows that they're actively working on it as of yesterday,
and it sounds as if they're some ways away from a solution. Once they
figure it out, it may be some time before the fix makes it into the
official Linux kernel, and longer still until it gets into a OpenVZ kernel.

If you're okay with applying the OpenVZ patches to your own kernel
source and building your own OpenVZ-capable kernel, your best bet may be
to sit tight and keep watching that bugzilla page until they fix this
bug in the kernel. Then grab that newest source and give it a try...

Sorry I don't have anything more constructive to say, but that seems to
be the status at the moment. :(

--

Gregor Mosheh / Greg Allensworth
System Administrator, HostGIS cartographic development & hosting services
<http://www.HostGIS.com/>

"Remember that no one cares if you can back up,
only if you can restore." - AMANDA
Send instant messages to your online friends <http://uk.messenger.yahoo.com>

Subject: Re: SATA HDD Problem
Posted by [vaverin](#) on Tue, 17 Jul 2007 02:50:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

Markus, Kirill

About http://bugzilla.kernel.org/show_bug.cgi?id=8650
We have not one but at least 3 different problems here:
1) interrupts-related issue on VIA hardware.

Comment #27 From Tejun Heo 2007-07-10 09:09:09:
2. The PATA and SATA controllers share a PCI IRQ line. The PATA controller also seem to be hardwired to 14 or 14/15 depending on controller mode. The driver/ide drivers use IRQ auto-detection and detect 14 for the PATA part while libata honors pdev->irq and use 20. The end result is the same tho. One of the two hosts lose ability to assert IRQ and everything falls down. It's definitely related to IRQ routing and is really peculiar. Well, I wouldn't expect anything less from the vendor. :-) Does "acpi=noirq" make any difference?

VvS: I would note that using libata drivers (instead ide) for PATA controller makes the situation much better, but unfortunately do not closes this issue completely: with using libata driver I've reproduced this issue, however it was only once. Also I would add that I still cannot reproduce this issue with "acpi=noirq".

2) Another issue is infinite Error Handling resets for ide-attached DVD-ROM. It is unpleasantly too because of it generates tonns of garbage messages in the system logs, however it brokes nothing on my node and therefore it have low severity for me.

3) ext3/jbd-related issue:
AIM7 test leads to the ext3/jbd lockup on 2.6.22-rc4 and -rc7 kenrels. However it looks like this issue is go away: I've updated the kernel up to 2.6.22 and still cannot reproduce it since Jul 12.

I know nothing about 2.6.9-based kernels. IMHO interrupt-related issue should be present on this kernels too, however I never saw bugreports until we have upgraded to 2.6.18 kernels.

Also I would note that all 3 issues are not Virtuozzo-specific and any new bugreports should be addressed to libata or ext3 developers but not to me, I'm just a tester in this situation.

Markus, your situation is not clear for me. I even cannot confirm that you have the same issue as I've observed. At the first glance all issues looks similar but can have the different reasons.

I do not know all details of your situation and may be wrong. However IMHO "device error" messages in your logs points to some disk drive failure. I would note that in my case this message was "timeout" and from my point of view it is important difference.

I would like to recommend you find the way to reproduce issue on your node, collect all information described your situation (all kernel messages beginning from node booting, lspci -vvvxxx output, probably something else) and send bugreport to libata developers.

If you don't want to investigate this bug -- I recommend you try to replace your hardware, beginning at disk drive.

Thank you,
Vasily Averin
