
Subject: Intel ahci errors on RHEL kernel
Posted by [lazy](#) on Tue, 29 Jan 2008 10:11:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi all,

After upgrading our system to ICH9 running in ahci mode.

Kernel starts to report sata exceptions. The disc is brand new (tested for 5 days didn't show any of these) after going to production it started to report these errors:

```
ata5.00: exception Emask 0x2 SAct 0x9 SErr 0x0 action 0x2 frozen
ata5.00: (spurious completions during NCQ issue=0x0 SAct=0x9 FIS=004040a1:00000004)
ata5.00: cmd 61/08:00:17:58:48/00:00:0d:00:00/40 tag 0 cdb 0x0 data 4096 out
         res 40/00:10:87:e5:25/00:00:0d:00:00/40 Emask 0x2 (HSM violation)
ata5.00: cmd 61/08:18:ef:e7:37/00:00:0d:00:00/40 tag 3 cdb 0x0 data 4096 out
         res 40/00:10:87:e5:25/00:00:0d:00:00/40 Emask 0x2 (HSM violation)
ata5: soft resetting port
ata5: SATA link up 3.0 Gbps (SStatus 123 SControl 300)
ata5.00: configured for UDMA/133
ata5: EH complete
SCSI device sda: 488397168 512-byte hdwr sectors (250059 MB)
sda: Write Protect is off
sda: Mode Sense: 00 3a 00 00
SCSI device sda: drive cache: write back
```

```
ata5.00: exception Emask 0x2 SAct 0x8 SErr 0x0 action 0x2 frozen
ata5.00: (spurious completions during NCQ issue=0x0 SAct=0x8 FIS=004040a1:00000004)
ata5.00: cmd 61/08:18:e7:b0:fb/00:00:17:00:00/40 tag 3 cdb 0x0 data 4096 out
         res 40/00:18:e7:b0:fb/00:00:17:00:00/40 Emask 0x2 (HSM violation)
ata5: soft resetting port
ata5: SATA link up 3.0 Gbps (SStatus 123 SControl 300)
ata5.00: configured for UDMA/133
ata5: EH complete
SCSI device sda: 488397168 512-byte hdwr sectors (250059 MB)
sda: Write Protect is off
sda: Mode Sense: 00 3a 00 00
SCSI device sda: drive cache: write back
```

here are the dmesg logs

```
ahci 0000:00:1f.2: AHCI 0001.0200 32 slots 6 ports 3 Gbps 0x3f impl SATA mode
ahci 0000:00:1f.2: flags: 64bit ncq led clo pmp pio slum part
PCI: Setting latency timer of device 0000:00:1f.2 to 64
scsi0 : ahci
scsi1 : ahci
scsi2 : ahci
scsi3 : ahci
scsi4 : ahci
scsi5 : ahci
```

ata1: SATA max UDMA/133 cmd 0xf8a7c100 ctl 0x00000000 bmdma 0x00000000 irq 74
ata2: SATA max UDMA/133 cmd 0xf8a7c180 ctl 0x00000000 bmdma 0x00000000 irq 74
ata3: SATA max UDMA/133 cmd 0xf8a7c200 ctl 0x00000000 bmdma 0x00000000 irq 74
ata4: SATA max UDMA/133 cmd 0xf8a7c280 ctl 0x00000000 bmdma 0x00000000 irq 74
ata5: SATA max UDMA/133 cmd 0xf8a7c300 ctl 0x00000000 bmdma 0x00000000 irq 74
ata6: SATA max UDMA/133 cmd 0xf8a7c380 ctl 0x00000000 bmdma 0x00000000 irq 74
ata1: SATA link down (SStatus 0 SControl 300)
ata2: SATA link down (SStatus 0 SControl 300)
ata3: SATA link down (SStatus 0 SControl 300)
ata4: SATA link down (SStatus 0 SControl 300)
ata5: SATA link up 3.0 Gbps (SStatus 123 SControl 300)
ata5.00: ATA-7: ST3250820NS, 3.AEG, max UDMA/133
ata5.00: 488397168 sectors, multi 0: LBA48 NCQ (depth 31/32)
ata5.00: configured for UDMA/133
ata6: SATA link down (SStatus 0 SControl 300)
Vendor: ATA Model: ST3250820NS Rev: 3.AE
Type: Direct-Access ANSI SCSI revision: 05
ACPI: PCI Interrupt 0000:02:00.0[A] -> GSI 17 (level, low) -> IRQ 169
PCI: Setting latency timer of device 0000:02:00.0 to 64
scsi6 : pata_marvell
scsi7 : pata_marvell
ata7: PATA max UDMA/100 cmd 0x00011018 ctl 0x00011026 bmdma 0x00011000 irq 169
ata8: DUMMY
BAR5:00:00 01:7F 02:22 03:CA 04:00 05:00 06:00 07:00 08:00 09:00 0A:00 0B:00 0C:01 0D:00
0E:00 0F:00
SCSI device sda: 488397168 512-byte hdwr sectors (250059 MB)
sda: Write Protect is off
sda: Mode Sense: 00 3a 00 00
SCSI device sda: drive cache: write back
SCSI device sda: 488397168 512-byte hdwr sectors (250059 MB)
sda: Write Protect is off
sda: Mode Sense: 00 3a 00 00
SCSI device sda: drive cache: write back
sda: sda1 sda2 sda3
sd 4:0:0:0: Attached scsi disk sda

the kernel is 2.6.18-53.el5.028stab051.1, with pae config the motherboard is an Intel with G33 express chipset

while googling i found this thread on lkm (
<http://www.mail-archive.com/linux-ide@vger.kernel.org/msg13658.html>)
so maybe it's just ahci driver fault but I'm not so fast to testing it on a production system

I'm thinking about disabling sata2 on that disc and hopefully disabling NCQ and mitigating this problem. Anyone has come against something similar ?

Any pointers will be appreciated.

Subject: Re: Intel achci errors on RHEL kernel
Posted by [piavlo](#) on Wed, 30 Jan 2008 22:41:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

AFAIK ICH9 is not supported by 2.6.18 vanilla kernel
From grepping kernel changelogs it looks like
ICH9 May have support starting from 2.6.20,
while 2.6.19 has a couple of ICH9 introducing patches, but not full support. While 2.6.18 and
earlier has no mention of ICH9 at all.

This makes openvz not usable for new ICH9 , while this is 90%
of intel based boards that on the market.
You could use 2.6.22 but it does not have checkpointing support &
live migration.

The same problem is with vanilla Xen 3.1 & 3.2 which also still in
2.6.18, but many distros have ported xen patches to newer kernels.

Subject: Re: Intel achci errors on RHEL kernel
Posted by [lazy](#) on Thu, 31 Jan 2008 09:37:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

piavlo wrote on Wed, 30 January 2008 17:41AFAIK ICH9 is not supported by 2.6.18 vanilla kernel
From grepping kernel changelogs it looks like
ICH9 May have support starting from 2.6.20,
while 2.6.19 has a couple of ICH9 introducing patches, but not full support. While 2.6.18 and
earlier has no mention of ICH9 at all.

This makes openvz not usable for new ICH9 , while this is 90%
of intel based boards that on the market.
You could use 2.6.22 but it does not have checkpointing support &
live migration.
kernels.

Thank You for you answer.

ICH9 is only supported in rhel based patches (it has backported drivers)

i resolved the issue by disabling NCQ in the kernel

this bug was fixed in December (can't get the url now) by removing spurious NCQ... detection in
ahci driver, the patch says that this detection was useful only on old hardware and that it's normal
behavior on modern ahci hardware.

I didn't have the guts to backport new driver myself to 2.6.18-rhel and test it on production system,
and we have none of the ICH9 for testing, i guess redhat will update their patchset for 2.6.18
(based on some grepping this didn't get into recent rhel openvz patch)

Subject: Re: Intel achci errors on RHEL kernel
Posted by [piavlo](#) on Thu, 31 Jan 2008 12:20:01 GMT

By the way i have AHCI disabled in the bios since then i enable it , then system powers on the first thing it does is staring the AHCI bios where it gets stuck for 5-10 minutes and then it continues normally. Do you have an idea what might be wrong, maybe i to change something else in the bios?

Regarding the backported RH patch , where can i get this separate backport only patch?

Thanks

Subject: Re: Intel achci errors on RHEL kernel
Posted by [lazy](#) on Thu, 31 Jan 2008 12:28:05 GMT

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

piavlo wrote on Thu, 31 January 2008 07:20 By the way i have AHCI disabled in the bios since then i enable it , then system powers on the first thing it does is staring the AHCI bios where it gets stuck for 5-10 minutes and then it continues normally. Do you have an idea what might be wrong, maybe i to change something else in the bios?

Regarding the backported RH patch , where can i get this separate backport only patch?

Thanks

we didn't experience it, i would try to upgrade bios, change cables, connect disks to other ports etc.

thers no separate patch for ahci, you have to go for full rhel package or do it yourself
