
Subject: unexpected scsi timeout
Posted by [vaverin](#) on Fri, 20 Jul 2007 12:02:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

Tejun, Jeff

I've noticed that some scsi commands for DVD-drive attached to pata_via successfully finishes without any delays but reports about TIMEOUT condition. It happens because of ATA_ERR bit is set in status register. As result for each command Error Handler thread awakened, requests sense buffer and go to sleep again.

Could you please confirm that it is expected behavior? It's a bit strange for me, IMHO other scsi drivers requests sense buffer without EH thread.

Thank you,
Vasily Averin

Subject: Re: unexpected scsi timeout
Posted by [Tejun Heo](#) on Tue, 24 Jul 2007 08:20:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

[cc'ing Albert]

Vasily Averin wrote:

> Tejun, Jeff
>
> I've noticed that some scsi commands for DVD-drive attached to pata_via
> successfully finishes without any delays but reports about TIMEOUT condition. It
> happens because of ATA_ERR bit is set in status register. As result for each
> command Error Handler thread awakened, requests sense buffer and go to sleep again.

Need more info. Please post boot dmesg and the result of 'lspci -nn' and 'hdparm -l /dev/srX' and when such errors occur.

Thanks.

--
tejun

Subject: Re: unexpected scsi timeout
Posted by [vaverin](#) on Tue, 24 Jul 2007 10:07:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

Tejun Heo wrote:
> [cc'ing Albert]

```
>
> Vasily Averin wrote:
>> Tejun, Jeff
>>
>> I've noticed that some scsi commands for DVD-drive attached to pata_via
>> successfully finishes without any delays but reports about TIMEOUT condition. It
>> happens because of ATA_ERR bit is set in status register. As result for each
>> command Error Handler thread awakened, requests sense buffer and go to sleep again.
>
> Need more info. Please post boot dmesg and the result of 'lspci -nn'
> and 'hdparm -I /dev/srX' and when such errors occur.
```

It was 2.6.22 kernel with pata_via and sata_via drivers, scsi and cdrom debug was temporally enabled via sysctl (please see logs near Jul 24 13:42:46 timestamp)

Btw. I'm not sure that it was an error, I've looked on the sources and IMHO it's normal command processing cdrom without disk inserted into drive. I've checked 2.6.19, 2.6.20 and 2.6.22 kernels and got the same behavior.

```
# lspci -nn
00:00.0 0600: 1106:3188 (rev 01)
00:01.0 0604: 1106:b188
00:0b.0 0200: 14e4:1653 (rev 03)
00:0f.0 0104: 1106:3149 (rev 80)
00:0f.1 0101: 1106:0571 (rev 06)
00:10.0 0c03: 1106:3038 (rev 81)
00:10.1 0c03: 1106:3038 (rev 81)
00:10.2 0c03: 1106:3038 (rev 81)
00:10.4 0c03: 1106:3104 (rev 86)
00:11.0 0601: 1106:3227
00:18.0 0600: 1022:1100
00:18.1 0600: 1022:1101
00:18.2 0600: 1022:1102
00:18.3 0600: 1022:1103
00:19.0 0600: 1022:1100
00:19.1 0600: 1022:1101
00:19.2 0600: 1022:1102
00:19.3 0600: 1022:1103
```

```
# hdparm -I /dev/cdrom
```

```
/dev/cdrom:
HDIO_DRIVE_CMD(identify) failed: Input/output error
```

(Hmm, strange. please see logs near Jul 24 13:44:30 timestamp)

You can find more details about this node in
http://bugzilla.kernel.org/show_bug.cgi?id=8650

I have full logs from kernel compiled with extra ata debug -- but it is very big. Please tell me if you want to look it too.

thank you,
Vasily Averin

Subject: Re: unexpected scsi timeout
Posted by [Albert Lee](#) on Wed, 25 Jul 2007 05:50:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Vasily Averin wrote:

> Tejun Heo wrote:

>

>>[cc'ing Albert]

>>

>>Vasily Averin wrote:

>>

>>>Tejun, Jeff

>>>

>>>I've noticed that some scsi commands for DVD-drive attached to pata_via
>>>successfully finishes without any delays but reports about TIMEOUT condition. It
>>>happens because of ATA_ERR bit is set in status register. As result for each
>>>command Error Handler thread awakened, requests sense buffer and go to sleep again.
>>

>>Need more info. Please post boot dmesg and the result of 'lspci -nn'

>>and 'hdparm -l /dev/srX' and when such errors occur.

>

>

> It was 2.6.22 kernel with pata_via and sata_via drivers, scsi and cdrom debug
> was temporally enabled via sysctl (please see logs near Jul 24 13:42:46 timestamp)

>

> Btw. I'm not sure that it was an error, I've looked on the sources and IMHO it's
> normal command processing cdrom without disk inserted into drive. I've checked
> 2.6.19, 2.6.20 and 2.6.22 kernels and got the same behavior.

>

Hi Vasily,

Your log looks ok. It's normal for TEST_UNIT_READY to return ATA_ERR when no disc
inside and libata EH triggered to request sense.

--

albert

Subject: Re: unexpected scsi timeout

Posted by [vaverin](#) on Wed, 25 Jul 2007 07:36:46 GMT

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

Albert Lee wrote:

>>> Vasily Averin wrote:

>>>> I've noticed that some scsi commands for DVD-drive attached to pata_via
>>>> successfully finishes without any delays but reports about TIMEOUT condition. It
>>>> happens because of ATA_ERR bit is set in status register. As result for each
>>>> command Error Handler thread awakened, requests sense buffer and go to sleep again.
>>> Need more info. Please post boot dmesg and the result of 'lspci -nn'
>>> and 'hdparm -l /dev/srX' and when such errors occur.

>
> Your log looks ok. It's normal for TEST_UNIT_READY to return ATA_ERR when no disc
> inside and libata EH triggered to request sense.

It's a bit strange for me, IMHO other scsi drivers requests sense buffer without
EH thread assistance.

Currently we know that ATA_ERR can be returned; it is not error, but one of
expected responses. Why we cannot request sense without EH? I would like to
understand is it implementation drawback or I missed something probably?

Thank you,
Vasily Averin

Subject: Re: unexpected scsi timeout

Posted by [Tejun Heo](#) on Wed, 25 Jul 2007 07:42:37 GMT

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

Vasily Averin wrote:

> Albert Lee wrote:

>>>> Vasily Averin wrote:

>>>>> I've noticed that some scsi commands for DVD-drive attached to pata_via
>>>>> successfully finishes without any delays but reports about TIMEOUT condition. It
>>>>> happens because of ATA_ERR bit is set in status register. As result for each
>>>>> command Error Handler thread awakened, requests sense buffer and go to sleep again.
>>>> Need more info. Please post boot dmesg and the result of 'lspci -nn'
>>>> and 'hdparm -l /dev/srX' and when such errors occur.

>> Your log looks ok. It's normal for TEST_UNIT_READY to return ATA_ERR when no disc
>> inside and libata EH triggered to request sense.

>
> It's a bit strange for me, IMHO other scsi drivers requests sense buffer without
> EH thread assistance.
> Currently we know that ATA_ERR can be returned; it is not error, but one of
> expected responses. Why we cannot request sense without EH? I would like to
> understand is it implementation drawback or I missed something probably?

That was a design choice. It's easier to implement that way.

--
tejun

Subject: Re: unexpected scsi timeout

Posted by [James Bottomley](#) on Wed, 25 Jul 2007 13:06:41 GMT

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

On Wed, 2007-07-25 at 16:42 +0900, Tejun Heo wrote:

> Vasily Averin wrote:

> > Albert Lee wrote:

> >>> Vasily Averin wrote:

> >>>> I've noticed that some scsi commands for DVD-drive attached to pata_via

> >>>> successfully finishes without any delays but reports about TIMEOUT condition. It

> >>>> happens because of ATA_ERR bit is set in status register. As result for each

> >>>> command Error Handler thread awakened, requests sense buffer and go to sleep again.

> >>> Need more info. Please post boot dmesg and the result of 'lspci -nn'

> >>> and 'hdparm -I /dev/srX' and when such errors occur.

> >> Your log looks ok. It's normal for TEST_UNIT_READY to return ATA_ERR when no disc

> >> inside and libata EH triggered to request sense.

> >

> > It's a bit strange for me, IMHO other scsi drivers requests sense buffer without

> > EH thread assistance.

> > Currently we know that ATA_ERR can be returned; it is not error, but one of

> > expected responses. Why we cannot request sense without EH? I would like to

> > understand is it implementation drawback or I missed something probably?

>

> That was a design choice. It's easier to implement that way.

And just so we're clear what SCSI allows:

On ordinary SCSI devices, when the device goes into a check condition state, it won't accept any more commands until it sees a request sense.

For SCSI devices this can be a problem (because there are several thousand sense conditions, some of which correspond to everything's alright), so a large number of SCSI drivers implement auto request sense emulation, which means that in the driver, as soon as they see the check condition, they immediately send a REQUEST SENSE command to pick up the sense code (minimising the time the device is blocked).

For drivers that don't want to implement this (and we have a few in SCSI) the alternative mechanism is to have the eh thread collect the sense data. This is the route libata has chosen.

James
