
Subject: irq0 stops working

Posted by [vaverin](#) on Tue, 09 Oct 2007 05:26:49 GMT

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

On one of our servers timer interrupts (i.e irq0) are stops working. As result any kernel timers do not triggers and tasks waiting some signals from timers hangs forever.

Most noticeable effect of this situation is that any write operations to disk are stalled, and nobody can log in on the node.

At the same time node all existing shells works away. I'm able to read interrupts statistic from /proc/interrupts file and it shows that all other interrupts are changed when these devices are accessed: disk on sata controller, network, cdrom on ide controller, keyboard, serial console, LOC interrupts.

Also I've found that disable of irqbalance service on the node helps to workaround this issue, however of course it fixes nothing.

All details about hardware/logs could be found in http://bugzilla.kernel.org/show_bug.cgi?id=8650

I'm able to reproduce this situation, however now I have no ideas how to continue the investigation of this problem.

Could please anybody advise me any new ways for investigation of this issue?

Thank you,
Vasily Averin

OpenVZ Linux Kernel Team

Subject: Re: irq0 stops working

Posted by [Jan Engelhardt](#) on Tue, 09 Oct 2007 05:34:08 GMT

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

On Oct 9 2007 09:26, Vasily Averin wrote:

>

>On one of our servers timer interrupts (i.e irq0) are stops working. As result
>any kernel timers do not triggers and tasks waiting some signals from timers
>hangs forever.

What kernel.. and tried CONFIG_NO_HZ=n?

Subject: Re: irq0 stops working

On Tue, 9 Oct 2007, Vasily Averin wrote:

> On one of our servers timer interrupts (i.e irq0) are stops working. As result
> any kernel timers do not triggers and tasks waiting some signals from timers
> hangs forever.

Which kernel version ?

> Most noticeable effect of this situation is that any write operations to disk
> are stalled, and nobody can log in on the node.
>
> At the same time node all existing shells works away. I'm able to read
> interrupts statistic from /proc/interrupts file and it shows that all other
> interrupts are changed when these devices are accessed: disk on sata controller,
> network, cdrom on ide controller, keyboard, serial console, LOC interrupts.
>
> Also I've found that disable of irqbalance service on the node helps to
> workaround this issue, however of course it fixes nothing.

Well, it's at least a hint. Can you try the patch below please ?

tgix

```
diff --git a/arch/x86_64/kernel/time.c b/arch/x86_64/kernel/time.c
index 6d48a4e..248987a 100644
--- a/arch/x86_64/kernel/time.c
+++ b/arch/x86_64/kernel/time.c
@@ -360,7 +360,7 @@ void stop_timer_interrupt(void)

static struct irqaction irq0 = {
    .handler = timer_interrupt,
-   .flags = IRQF_DISABLED | IRQF_IRQPOLL,
+   .flags = IRQF_DISABLED | IRQF_IRQPOLL | IRQF_NOBALANCING,
    .mask = CPU_MASK_NONE,
    .name = "timer"
};
@@ -403,6 +403,7 @@ void __init time_init(void)
    cpu_khz / 1000, cpu_khz % 1000);
    init_tsc_clocksource();

+   irq0.mask = cpumask_of_cpu(0);
    setup_irq(0, &irq0);
}
```

Subject: Re: irq0 stops working

Posted by [vaverin](#) on Tue, 09 Oct 2007 06:12:15 GMT

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

Jan Engelhardt wrote:

> On Oct 9 2007 09:26, Vasily Averin wrote:

>> On one of our servers timer interrupts (i.e irq0) are stops working. As result
>> any kernel timers do not triggers and tasks waiting some signals from timers
>> hangs forever.

>

> What kernel.. and tried CONFIG_NO_HZ=n?

Originally I've noticed this issue on RHEL5 kernels, but then I've reproduced it on latest mainstream kernels, in my last attempt it was 2.6.23-rc7.

Thank you for for your tips about CONFIG_NO_HZ=n, will try to to it.

thank you,
Vasily Averin

OpenVZ Linux Kernel Team

Subject: Re: irq0 stops working

Posted by [Thomas Gleixner](#) on Tue, 09 Oct 2007 06:23:33 GMT

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

On Tue, 9 Oct 2007, Vasily Averin wrote:

> Jan Engelhardt wrote:

> > On Oct 9 2007 09:26, Vasily Averin wrote:

> >> On one of our servers timer interrupts (i.e irq0) are stops working. As result
> >> any kernel timers do not triggers and tasks waiting some signals from timers
> >> hangs forever.

> >

> > What kernel.. and tried CONFIG_NO_HZ=n?

>

> Originally I've noticed this issue on RHEL5 kernels, but then I've reproduced it
> on latest mainstream kernels, in my last attempt it was 2.6.23-rc7.

>

> Thank you for for your tips about CONFIG_NO_HZ=n, will try to to it.

You run a 64 bit kernel, where this option is not available yet.

tglx

Subject: Re: irq0 stops working

Posted by [vaverin](#) on Tue, 06 Nov 2007 07:15:29 GMT

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

Thomas Gleixner wrote:

> On Tue, 9 Oct 2007, Vasily Averin wrote:

>> On one of our servers timer interrupts (i.e irq0) are stops working. As result
>> any kernel timers do not triggers and tasks waiting some signals from timers
>> hangs forever.

>>

>> Also I've found that disable of irqbalance service on the node helps to
>> workaround this issue, however of course it fixes nothing.

>

> Well, it's at least a hint. Can you try the patch below please ?

Patch helps, after 1-month-long testing I could not reproduce this issue again,
what are we going to do now?

Unfortunately my test node is still unstable, now it reproduces sata timeout
issue, details are here:

http://bugzilla.kernel.org/show_bug.cgi?id=8650

Thank you,
Vasily Averin
