

Re: pcim_enable_device BUGs for libata devices in 2.6.20-gi6

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-02/msg04187.html>

- *From:* Pavel Machek <pavel@xxxxxx>
 - *Date:* Mon, 12 Feb 2007 21:54:44 +0100
-

Hi!

I'm seeing BUGs like these on all libata-driven controllers when suspending to disk on 2.6.20-gi6:

```
sata_nv 0000:00:07.0: resuming
BUG: at drivers/pci/pci.c:817 pcim_enable_device()
```

Call Trace:

```
[<ffffffff80337d21>] pcim_enable_device+0x8a/0xa5
[<ffffffff88099d18>] :libata:ata_pci_device_do_resume+0x20/0x59
[<ffffffff880bb731>] :sata_nv:nv_pci_device_resume+0x1d/0x100
[<ffffffff8039d2bf>] resume_device+0xcb/0x12c
[<ffffffff8039d3ac>] dpm_resume+0x8c/0xec
[<ffffffff8039d456>] device_resume+0x4a/0x5d
[<ffffffff802a0a33>] pm_suspend_disk+0x160/0x170
[<ffffffff8029f4b6>] enter_state+0x52/0x1da
[<ffffffff8029f69c>] state_store+0x5e/0x79
[<ffffffff802f2b20>] sysfs_write_file+0xe4/0x118
[<ffffffff80214b58>] vfs_write+0xce/0x177
[<ffffffff8021553e>] sys_write+0x45/0x6e
[<ffffffff8025711e>] system_call+0x7e/0x83
```

It looks like what's happening is that during the "freezing" stage, we suspend and then resume the controllers. ata_pci_device_do_suspend only calls pci_disable_device if the event is PM_EVENT_SUSPEND but ata_pci_device_do_resume calls pcim_enable_device unconditionally. If the event was something else, then pcim_enable_device complains because the device was previously enabled and never disabled.

Not sure what the best way to fix this is?

I think what should happen is either one of the followings.

1. Don't restore power state and re-enable PCI device on resume from freeze just as we don't do the opposite when freezing.

2. Unconditionally disable and power down PCI device on suspend whether it's freeze or not.

#2 would be simpler but I'm a bit worried about it. There might be controllers which choke after such sequence (save state, disable, power down, no actual power removal, power on, restore state, re-enable).

I'd just go for #2.

#1 can be easily done by taking a look at the current device power state (gendev->power.power_state). The problem is that no one in suspend/resume path seems to be setting that variable except for runtime

No, that variable is probably going to go away. If you want to remember that you are resuming from freeze, just store that info in private data structure.

Pavel

--

(english) <http://www.livejournal.com/~pavelmachek>

(cesky, pictures) <http://atrey.karlin.mff.cuni.cz/~pavel/picture/horses/blog.html>

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>