Subject: Re: [PATCH] pci_get_device call from interrupt in reboot fixups Posted by Andrew Morton on Tue, 07 Aug 2007 07:44:55 GMT

View Forum Message <> Reply to Message

On Tue, 7 Aug 2007 00:24:37 -0700 Andrew Morton <akpm@linux-foundation.org> wrote:

```
>> Andrew, I really don't want to change the PCI core to handle this, as we
> > finally fixed a lot of issues with drivers trying to walk these lists
>> from interrupt context. So if you want to just hide the warning message
>> as we are shutting down, that's fine with me. Or just don't do the
>> fixups. But grabbing a reference to the pci device is unsafe in my
> > opinion and I do not want to do that.
> >
>
> OK, good decision ;)
> One approach would be for some brave soul to pick his way through
> the reboot code and ensure that we are correctly and reliably setting
> system state to SYSTEM RESTART, then test that in might sleep().
> But this does suppress somewhat-useful debugging just because of sysrq-B
> and I really wouldn't want to utilise the horrid system state any more that
> we are presently doing. I think on balance that it would be better if we
> could do something more targetted, like modify emergency_restart() to test
> in_interrupt() and to then apologetically set some well-named global flag
> which will shut up might sleep(). Pretty foul, but I can't think of
> anything better.
ok, this might be better. How about we just stop calling mach reboot fixups()
at sysrq-B time?
>>> handle_sysrq
        machine_emergency_restart
> > >>
          mach_reboot_fixups
> > >>
           pci_get_device
> > >>
            pci get subsys
         down_read
> > >>
```