

Re: [PATCH] x86: Document the hotplug code is incompatible with x86 irq handling

## Re: [PATCH] x86: Document the hotplug code is incompatible with x86 irq handling

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-05/msg13796.html>

---

- *From:* [ebiederm@xxxxxxxxxxxxx](mailto:ebiederm@xxxxxxxxxxxxx) (Eric W. Biederman)
  - *Date:* Thu, 31 May 2007 09:47:22 -0600
- 

Robert Hancock <[hancockr@xxxxxxx](mailto:hancockr@xxxxxxx)> writes:

Eric W. Biederman wrote:

I just realized that except for doing the code review and noticing that the current cpu hotplug code is fundamentally incompatible with x86 I haven't done anything about it. So here is my patch to document what is wrong.

The current cpu hotplug code requires irqs to be migrated from a cpu outside of irq context. On x86 ioapics simply do not support this, making the code unfixable without major redesign of the generic cpu hotplug code.

So this patch makes CPU\_HOTPLUG on x86 depend on CONFIG\_BROKEN and adds a WARN\_ON so people that do enable it are not in doubt about which part of the code is broken, even if it does work for them.

Signed-off-by: Eric W. Biederman <[ebiederm@xxxxxxxxxxxxx](mailto:ebiederm@xxxxxxxxxxxxx)>

I don't think this is useful, though the code may be problematic, this patch will break suspend on all SMP machines with an existing config, which is a major regression..

Cool.

My patch does not change the functionality of the code just complains very loudly that it is broken.

Further the code is broken at a design level. The code isn't problematic the code is impossible. The cpu hotplug code can not be fixed on x86 without a redesign of the generic cpu hotplug code.

Re: [PATCH] x86: Document the hotplug code is incompatible with x86 irq handling

Suspend does not need to use cpu hotplug because it already gets in deep with the drivers, and can stop interrupts at the source. I know there was some talk about this doing this earlier, but I don't know if anything came of that discussion.

Regardless if you care about this being a problem feel free to fix the relevant code so it attempts to do something that the hardware actually supports.

But if the suspend needs this code for smp support it is also broken.

Eric

–

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

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

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