

Re: ANNOUNCE: CE Linux Forum – Specification V1.0 draft

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-05/3727.html>

From: Russell King (rmk+lkml_at_arm.linux.org.uk)

Date: 05/18/04

Date: Tue, 18 May 2004 20:56:08 +0100
To: Mark Gross <mgross@linux.jf.intel.com>

On Tue, May 18, 2004 at 12:32:48PM -0700, Mark Gross wrote:

```
> > --- linux-2.4.20.orig/drivers/ide/ide.c Thu Nov 28 23:53:13 2002
> > +++ celinux-040213/drivers/ide/ide.c Thu Feb 12 10:25:12 2004
> > @@ -2739,12 +2776,17 @@
> > */
> > void ide_delay_50ms (void)
> > {
> > +#ifdef CONFIG_IDE_PREEMPT
> > + __set_current_state(TASK_UNINTERRUPTIBLE);
> > + schedule_timeout(1+HZ/20); /* from 2.5 */
> > +#else /* CONFIG_IDE_PREEMPT */
> > #ifndef CONFIG_BLK_DEV_IDECS
> > mdelay(50);
> > #else
> > __set_current_state(TASK_UNINTERRUPTIBLE);
> > schedule_timeout(HZ/20);
> > #endif /* CONFIG_BLK_DEV_IDECS */
> > +#endif /* CONFIG_IDE_PREEMPT */
> > }
> >
> > This great piece 'called IDE-preempt' to be buzzword-compliant is (and
> > that's noticeable just from looking at the diff!) so braindead that it's
> > not explainable by incompetence alone. You'd get your same result by just
> > _disabling_ CONFIG_BLK_DEV_IDECS instead of adding another broken config
> > option (modulo 2.6 adjustments to the sleep time).
> >
> > Every engineer with the slightest clue would first disable that option, or
> > if ide-cs support is actually needed think _why_ it's different instead of
> > just adding a config option to disable it. Either it's safe to always use
> > the sleeping variant in which case the original ifdef should go away, or
> > it's not in which case your patch is completely broken.
> >
> >
> > OK I'll bite, but just because in your blind hostility and haste you've made a
> > mistake ;)
```

Christoph is actually making a valid point here and I suspect is trying to point out the lack of thought put into this change. The things that `_should_` have been considered before making the change are:

1. Why do we use `mdelay()` here if `CONFIG_BLK_DEV_IDECS` is defined?
2. Is the reason for this still valid?
3. If it is safe to sleep here even if `CONFIG_CLK_DEV_IDECS` is set, why bother with `mdelay()` in the first place?

The `_correct_` patch is actually:

```
void ide_delay_50ms (void)
{
-#ifndef CONFIG_BLK_DEV_IDECS
- mdelay(50);
-#else
    __set_current_state(TASK_UNINTERRUPTIBLE);
    schedule_timeout(HZ/20);
-#endif /* CONFIG_BLK_DEV_IDECS */
}
```

since PCMCIA always calls drivers from process context now.

(Unfortunately I can't write upside down, but I'll give the answers to those three items above. Look away now if you don't want to read the answers! 8))

1. PCMCIA used to call drivers in IRQ context, which made it impossible to sleep.
2. No, because PCMCIA always calls drivers in process context now, so sleeping is possible.
3. Left as an exercise to the reader. 8)

```
--
Russell King
Linux kernel    2.6 ARM Linux    - http://www.arm.linux.org.uk/
maintainer of: 2.6 PCMCIA      - http://pcmcia.arm.linux.org.uk/
                2.6 Serial core
```

```
-
To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
the body of a message to majordomo@vger.kernel.org
More majordomo info at http://vger.kernel.org/majordomo-info.html
Please read the FAQ at http://www.tux.org/lkml/
```