

Re: question on resume()

## Re: question on resume()

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-01/msg08424.html>

---

- *From:* Oliver Neukum <[oneukum@xxxxxxx](mailto:oneukum@xxxxxxx)>
  - *Date:* Wed, 31 Jan 2007 10:04:08 +0100 (MET)
- 

Am Mittwoch, 31. Januar 2007 09:49 schrieb Rafael J. Wysocki:

On Wednesday, 31 January 2007 09:40, Oliver Neukum wrote:

Am Mittwoch, 31. Januar 2007 09:33 schrieb Rafael J. Wysocki:

On Tuesday, 30 January 2007 23:32, Rafael J. Wysocki wrote:

Generally, you are safe if your driver only calls `wake_up()` from a process context, but not from `.resume()` or `.suspend()` routines (or from an unfreezeable kernel thread).

Ah, sorry, I've just realized I was wrong. Processes in `TASK_UNINTERRUPTIBLE` cannot be frozen! So, the above only applies to `wake_up_interruptible()`.

So the kernel will wait for tasks in `TASK_UNINTERRUPTIBLE` to finish IO before it calls `suspend()`? I am confused.

Yes, it will. The process freezer can only return success if there are no more `TASK_UNINTERRUPTIBLE` tasks. Otherwise it fails (after a timeout).

So, this means, on `suspend()`:

1. Don't worry about `TASK_UNINTERRUPTIBLE`
  2. Do worry about `TASK_INTERRUPTIBLE`
- We have to cease IO and must not call `wake_up_interruptible()`

Isn't that a race until `suspend()` is called?

Re: question on resume()

Re: question on resume()

On resume():

1. Don't worry about TASK\_UNINTERRUPTIBLE
2. Do not restart IO that may call wake\_up\_interruptible()

When do we restart such IO?

Regards

Oliver

—

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

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

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