

Re: Log flood: "scheduling while atomic" (2.6.15.x, 2.6.16.x)

Re: Log flood: "scheduling while atomic" (2.6.15.x, 2.6.16.x)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-04/msg06720.html>

- *From:* Darren Salt <linux@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Sun, 30 Apr 2006 15:21:55 +0100
-

I demand that Andrew Morton may or may not have written...

Darren Salt <linux@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote:

I'm seeing bouts of log flooding caused by something presumably not releasing a lock. I've looked at some of the messages, but at around 100/s, I'm not too keen to look through the whole lot :-)

```
scheduling while atomic: swapper/0xafbffff/0
[show_trace+19/32]
[dump_stack+30/32]
[schedule+1278/1472]
[cpu_idle+88/96]
[stext+44/64]
[start_kernel+574/704]
[L6+0/2] 0xc0100199
```

(Trailing parts of some lines have been omitted; it's all repeated data. And some sort of rate-limiting of these messages would be nice, but some other way to draw attention to the problem, e.g. an occasional beep, would be good.)

The most recent instance occurred a few minutes into recording a TV programme (via vdr) from a cx88-based Nova-T. (I'm currently using stock drivers rather than ones built from the v4l-dvb repository.)

Thanks for the report.

Re: Log flood: "scheduling while atomic" (2.6.15.x, 2.6.16.x)

Re: Log flood: "scheduling while atomic" (2.6.15.x, 2.6.16.x)

The below patch (against 2.6.17-rc3) should, if it still works, tell us which lock didn't get unlocked.

You'll need to enable CONFIG_PREEMPT and CONFIG_DEBUG_PREEMPT and CONFIG_FRAME_POINTER.

Done, compiled, rebooted. I have a recording scheduled for later; I'll wait and see what happens.

Please cc video4linux-list@xxxxxxxxxxx on any result if it looks like v4l is indeed the culprit.

Will do.

BTW, patches applied:

- * the advansys patch from -mm;
- * BROKEN removed from the depends for advansys;
- * quietening of dprintk(0,...) (replaced with dprintk(1,...) in cx88-mpeg.c (these messages have some annoyance value);
- * a patch of my own for usbhid for a slightly weird USB+PS/2 mouse, connected via USB (I'll post this as directed in MAINTAINERS);
- * another of my own (attached for reference) which *should* rate-limit the "scheduling while atomic" messages somewhat.

The last one is new; the rest don't have any bearing on the problem, which has occurred without them and, indeed, without the presence of advansys and usbhid.

[snip]

--

| Darren Salt | linux or ds at | nr. Ashington, | Toon
| RISC OS, Linux | youmustbejoking,demon,co,uk | Northumberland | Army
| + Lobby friends, family, business, government. WE'RE KILLING THE PLANET.

He is truly wise who gains wisdom from another's mishap.

--- 2.6.17-rc3/kernel/sched.c.orig

+++ 2.6.17-rc3/kernel/sched.c

@@ -2904,10 +2904,28 @@

* Otherwise, whine if we are scheduling when we should not be.

*/

```
if (unlikely(in_atomic() && !current->exit_state)) {
```

```
+ /* Hack to avoid *serious* log-flooding. */
```

```
+ static int skipped = -50; /* want to report the first 50 */
```

```
+ static unsigned long last = 0;
```

Re: Log flood: "scheduling while atomic" (2.6.15.x, 2.6.16.x)

Re: Log flood: "scheduling while atomic" (2.6.15.x, 2.6.16.x)

```
+ int doprint = 1;
+ if (skipped < 0) {
+ if (!++skipped)
+ last = jiffies;
+ } else if (jiffies - last < 5 * HZ) /* should be 5s */ {
+ if (skipped < 0x7FFFFFFF)
+ ++skipped;
+ doprint = 0;
+ }
+ if (doprint) {
+ last = jiffies;
+ if (skipped)
+ printk(KERN_ERR "[%d s-w-a not reported]\n", skipped);
+ skipped = 0;
+ printk(KERN_ERR "BUG: scheduling while atomic: "
+ "%s/0x%08x/%d\n",
+ current->comm, preempt_count(), current->pid);
+ dump_stack();
+ }
+ }
+ }
profile_hit(SCHED_PROFILING, __builtin_return_address(0));
```