

[PATCH] Avoid moving tasks when a schedule can be made.

[PATCH] Avoid moving tasks when a schedule can be made.

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-01/msg11509.html>

- *From:* Steven Rostedt <rostedt@xxxxxxxxxxxx>
 - *Date:* Tue, 31 Jan 2006 14:43:29 -0500
-

I found this in the -rt kernel. While running "hackbench 20" I hit latencies of over 1.5 ms. That is huge! This latency was created by the move_tasks function in sched.c to rebalance the queues over CPUS.

There currently isn't any check in this function to see if it should stop, thus a large number of tasks can drive the latency high.

With the below patch, (tested on -rt with latency tracing), the latency caused by hackbench disappeared below my notice threshold (100 usecs).

I'm not convinced that this bail out is in the right location, but it worked where it is. Comments are welcome.

— Steve

Signed-off-by: Steven Rostedt <rostedt@xxxxxxxxxxxx>

Index: linux-2.6.16-rc1/kernel/sched.c

```
=====
--- linux-2.6.16-rc1.orig/kernel/sched.c 2006-01-19 15:58:52.000000000 -0500
+++ linux-2.6.16-rc1/kernel/sched.c 2006-01-31 14:27:17.000000000 -0500
@@ -1983,6 +1983,10 @@
```

```
curr = curr->prev;
```

```
+ /* bail if someone else woke up */
+ if (need_resched())
+ goto out;
+
if (!can_migrate_task(tmp, busiest, this_cpu, sd, idle, &pinned)) {
if (curr != head)
goto skip_queue;
```

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>

[PATCH] Avoid moving tasks when a schedule can be made.

[PATCH] Avoid moving tasks when a schedule can be made.

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

- Prev by Date: [*Re: GPL V3 and Linux – Dead Copyright Holders*](#)
- Next by Date: [*Re: GPL V3 and Linux – Dead Copyright Holders*](#)
- Previous by thread: [*\[PATCH\] Fix compile error in latest git pull – post 2.6.16-rc1-git4*](#)
- Next by thread: [*\[patch 0/6\] Add MMC password protection \(lock/unlock\) support V4*](#)
- Index(es):
 - ◆ [*Date*](#)
 - ◆ [*Thread*](#)