

Re: 2.6.21-mm1: many processes end up in D state

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-04/msg12181.html>

- *From:* Andrew Morton <akpm@xxxxxxxxxxxxxxxxxxxxxxxx>
 - *Date:* Mon, 30 Apr 2007 11:44:49 -0700
-

On Mon, 30 Apr 2007 20:14:05 +0200
"Jiri Slaby" <jirislaby@xxxxxxxxxx> wrote:

I have a problem with higher disk loads (e.g. running git-log or yum update).
Many processes end up in D state and system is unusable ---
I'm not able to run
anything but smooth mouse moving when this happens.

If I wait for a 20-30sec it becomes usable. This happens in
2.6.21-rc7-mm2 and
also in 2007-04-28-05-06 broken-out snapshot. I think
2.6.21-rc6-mm1 worked
fine, but I'm uncertain. If it is important, let me know to
re-test.

It is important, but I doubt if retesting 2.6.21-rc6-mm1 will clarify things a lot.

Could you try switching to a different IO scheduler please? Anticipatory would suit.

As I wrote below the sysrq-t, switch to noop didn't help, but it seems that it's harder to reproduce with that:

<cite it's_bad_to_write_anything_below_logs="true">
Note that yum works on lvm on raid0 and git too, but on the another md volume. Both ext3s. Drivers are sata Promise and ata_piix (sata disk); CFQ scheduler. Using noop is no change (but seems to be harder to reproduce with it). I figured out that it probably happens when 2+ processes are on both "processors" (HT on P4) and are IO wait (multiloader-applet shows red above the half).

Swap usage is 0 all the time.
</cite>

Re: 2.6.21-mm1: many processes end up in D state

My comprehension skills on Monday morning are even less than usual ;)

I would check the anticipatory scheduler as well, please. I don't know what no-op would do with a workload like that, but it probably isn't very good.

You appear to believe that it's related to the CPU scheduler? That's a bit unexpected – it sounds more like a VFS/IO thing? But stranger things have happened.

I guess it's time to end the staircase experiment in -mm.
<http://userweb.kernel.org/~akpm/js.bz2> is my current rollup (against 2.6.21) minus staircase and related things. Pretty please.

–

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