

Re: Xterm Hangs – Possible scheduler defect?

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-02/5787.html>

From: Mike Galbraith (efault_at_gmx.de)

Date: 02/25/05

Date: Fri, 25 Feb 2005 05:25:18 +0100
To: "Chad N. Tindel" <chad@tindel.net>

At 12:53 PM 2/24/2005 –0500, Chad N. Tindel wrote:

> > > *Hmmm... Are you suggesting it is OK for a kernel to get nearly completely
> > > hosed and for not fully utilize all the processors in the system because
> > > of one SCHED_FIFO thread?*

> >

> > *Sure. You specifically directed the scheduler to run your thread at a
> > higher priority than anything else. The way I see it, you used root's
> > perogative to shoot himself in the foot. You could also have used root's
> > perogative to don steel toed shoes(set important kernel threads to a higher
> > priority) before pulling the trigger.*

>

> *No, I specifically directed the scheduler to run my thread at a higher
> priority than any other userspace application. The fact that I wrote it
> in userspace and not in kernel space implies that I am OK with the kernel
> stopping me sometimes when it has work to do. If I wanted something
> higher priority than the kernel I would have written something in kernel
> space instead.*

Nope. You may have thought you told it that, but the reality is as I described it.

> > *SCHED_FIFO thread are supposed to preempt
> > > all other userspace threads... not the kernel itself.*

> >

> > *Not so. The scheduler makes do distinction between user and kernel threads
> > of execution.*

>

> *That is SOOOO broken it isn't even funny.*

I heartily disagree. I call it flexible/powerful.

> > *If you think that's broken, you'll love Ingo's IRQ threads...*

>

> *Yeah, thats broken too.*

(You're not noticing the added power it gives you.)

Linux-Kernel: Re: Xterm Hangs – Possible scheduler defect?

*>Perhaps I don't understand this philosophy you have where the kernel
>isn't more important than everything else. It seems to me like there needs
>to be a rigid hierarchy for scheduling, lest you get into deadlock problems:*

Some kernel thread flushing buffers should be more important than my
userland trigger-pacemaker thread?

*>Under no circumstances should any single CPU-bound userspace thread
>completely
>hose a 64-way SMP box.*

I can certainly agree that any service which is required across processor
borders wants to be very high priority indeed, and I can further agree that
this crossing of borders would not exist in a perfect world.

–Mike

–

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/>