

Re: [RFC][PATCH] O(1) Entitlement Based Scheduler

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-02/7774.html>

From: Joachim B Haga (*c.j.b.haga_at_fys.uio.no*)

Date: 02/29/04

To: Peter Williams <peterw@arema.com>

Date: 29 Feb 2004 12:58:14 +0100

Peter Williams <peterw@arema.com> writes:

>>> *They already do e.g. renice is such a program.*
>> *No one's talking about LOWERING priority here. You can only DoS*
>> *someone else if you can set negative nice values, and non-root*
>> *can't do that.*
>
> *Which is why root has to be in control of the mechanism.*

It seems to me that much of this could be solved if the user **were** allowed to lower nice values (down to 0).

Right now the only way I can prioritize between my own processes by starting important/timing sensitive programs normally and everything else reniced. The problem is that the first category consists of one or two programs while the second category is, well, "everything else".

I would **love** to be able to start the window manager and all children at +10 and be able to adjust priorities, from 0 (important user-level) to 10 (normal) to 20. Negative values could still be root-only.

So why shouldn't this be possible? Because a greedy user in a multi-user system would just run everything at max prio thus defeating the purpose? Sure, that would be annoying but it would have another solution ie. an entitlement based scheduler or something.

(and isn't it this simple?)

```
--- linux-2.6.3-mm3/kernel/sys.c.orig 2004-02-29 12:58:45.000000000 +0100
+++ linux-2.6.3-mm3/kernel/sys.c 2004-02-29 12:59:20.000000000 +0100
@@ -276,7 +276,7 @@
     error = -EPERM;
     goto out;
 }
- if (niceval < task_nice(p) && !capable(CAP_SYS_NICE)) {
```

Linux-Kernel: Re: [RFC][PATCH] O(1) Entitlement Based Scheduler

```
+ if (niceval < 0 && !capable(CAP_SYS_NICE)) {  
    error = -EACCES;  
    goto out;  
}
```

Regards,
Joachim B Haga

—

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/>