
Subject: Re: [PATCH, v6 2/3] Implement timer slack notifier chain
Posted by [Thomas Gleixner](#) on Mon, 14 Feb 2011 15:16:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 14 Feb 2011, Kirill A. Shutemov wrote:

> On Mon, Feb 14, 2011 at 02:32:23PM +0100, Thomas Gleixner wrote:
> > On Mon, 14 Feb 2011, Kirill A. Shutsemov wrote:
> >
> > > From: Kirill A. Shutemov <kirill@shutemov.name>
> > >
> > > Process can change its timer slack using prctl(). Timer slack notifier
> > > call chain allows to react on such change or forbid it.
> >
> > So we add a notifier call chain and more exports to allow what ?
>
> To allow the cgroup controller validate the value.

So we add 5 exports and a notifier chain to have a module? Errm, I mean there is not really a high probability that we'll add 5 more of those validation thingies, right?

```
So instead of having
#ifdef CONFIG_CGROUP_MUCK
int cgroup_set_slack(...);
#else
static inline int cgroup_set_slack(...)
{
    return ....
}
#endif
```

We add all that stuff ?

```
> > > --- a/kernel/sys.c
> > > +++ b/kernel/sys.c
> > > @@ -1691,15 +1691,10 @@ SYSCALL_DEFINE5(prctl, int, option, unsigned long, arg2,
> > > unsigned long, arg3,
> > >     error = perf_event_task_enable();
> > >     break;
> > >     case PR_GET_TIMERSLACK:
> > > -   error = current->timer_slack_ns;
> > > +   error = prctl_get_timer_slack();
> >
> > What's the point of replacing current->timer_slack_ns with a
> > function which does exactly the same ?
>
> To keep it consistent. BTW, prctl_get_seccomp() does the same.
```

That does not make it less bloat.

```
> >
> > > +long prctl_set_timer_slack(long timer_slack_ns)
> > > +{
> > > + int err;
> > > +
> > > + /* Reset timer slack to default value */
> > > + if (timer_slack_ns <= 0) {
> > > + current->timer_slack_ns = current->default_timer_slack_ns;
> > > + return 0;
> >
> > That does not make any sense at all. Why is setting
> > default_timer_slack_ns not subject to validation ?
>
> Hm.. In case of cgroup_timer_slack it's always valid.
> But, yes, in general, we should validate it.
>
> > Why is it treaded seperately ?
>
> What do you mean?
```

Should have read:

Why is it treated seperately from the other settings?

So setting the default is probably correct to be out of the validation thing, still the question remains, why we do not have a cgroup default then.

Thanks,

tgix

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
