
Subject: Re: [PATCH, v6 3/3] cgroups: introduce timer slack controller

Posted by [Matt Helsley](#) on Mon, 14 Feb 2011 23:39:36 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Tue, Feb 15, 2011 at 12:39:39AM +0200, Kirill A. Shutemov wrote:

> On Mon, Feb 14, 2011 at 06:01:06PM +0100, Thomas Gleixner wrote:

> > B1;2401;0cOn Mon, 14 Feb 2011, Kirill A. Shutemov wrote:

> >

> > > On Mon, Feb 14, 2011 at 03:00:03PM +0100, Thomas Gleixner wrote:

> > > > On Mon, 14 Feb 2011, Kirill A. Shutsemov wrote:

> > > > > From: Kirill A. Shutemov <kirill@shutemov.name>

<snip>

> > > > + list_for_each_entry(cur, &cgroup->children, sibling) {

> > > > + child = cgroup_to_tslack_cgroup(cur);

> > > > + if (type == TIMER_SLACK_MIN && val > child->min_slack_ns)

> > > > + return -EBUSY;

> > > >

> > > I thought the whole point is to propagate values through the group.

> > >

> > > I think silent change here is wrong. cpuset returns -EBUSY in similar

> > > case.

> >

> > And how is cpuset relevant for this ? Not at all. This is about

I agree with Thomas here -- cpusets aren't relevant.

> > timer_slack and we better have a well defined scheme for all of this

> > and not some cobbled together thing with tons of exceptions and corner

> > cases. Of course undocumented as far the code goes.

>

> I don't like silent cascade changes. Userspace can implement it if

It need not be totally silent. memcg has a "use_hierarchy" flag file.

Alternately, you could punt for now and disable hierarchy somewhat like blkio does.

> needed. -EBUSY is appropriate.

Hmm, I haven't thought about that method of cascading enough. The important question to consider is how will the parent cgroup be constrained if the owner/group of the children is different and thus disallows userspace from implementing this cascade. I suppose it's consistent with the owner/group ids but it hardly seems consistent with the "spirit" of using cgroups to enable things like containers.

Cheers,

-Matt Helsley

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
