
Subject: Re: containers (was Re: -mm merge plans for 2.6.23)
Posted by [Srivatsa Vaddagiri](#) on Wed, 11 Jul 2007 13:14:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, Jul 11, 2007 at 01:30:40PM +0200, Peter Zijlstra wrote:

> > One thing to think on though, we cannot have per process,uid,gid,pgrp
> > scheduling for one release only. So we'd have to manage interaction with
> > process containers. It might be that a simple weight multiplication
> > scheme is good enough:
> >
> > weight = uid_weight * pgrp_weight * container_weight

We would need something like this to flatten hierarchy, so that for example it is possible to do fair-container scheduling + fair-user/process scheduling inside a container using a hierarchy depth of just 1 (containers) that core scheduler understands. We discussed this a bit at <http://marc.info/?l=linux-kernel&m=118054481416140&w=2> and is very much on my todo list to experiment with.

> > Of course, if we'd only have a single level group scheduler (as was
> > proposed IIRC) it'd have to create intersection sets (as there might be
> > non trivial overlaps) based on these various weights and schedule these
> > resulting sets instead of the initial groupings.

>
> Lets illustrate with some ASCII art:
>
> so we have this dual level weight grouping (uid, container)
>
> uid: a a a a a b b b b b c c c c c
> container: A A A A A A B B B B B B B B
>
> set: 1 1 1 1 1 2 2 3 3 3 4 4 4 4 4
>
> resulting in schedule sets 1,2,3,4

Wouldn't it be simpler if admin created these sets as containers directly? i.e:

uid: a a a a a b b b b b c c c c c
container: 1 1 1 1 1 2 2 3 3 3 4 4 4 4 4

That way scheduler will not have to "guess" such intersecting schedulable sets/groups. It seems much simpler to me this way.

Surely there is some policy which is driving some tasks of userid 'b' to be in container A and some to be in B. It should be trivial enough to hook onto that policy making script and create separate

containers like above.

> so that (for instance) $\text{weight_2} = \text{weight_b} * \text{weight_A}$

--

Regards,
vatsa

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