
Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem

Posted by [Srivatsa Vaddagiri](#) on Wed, 04 Apr 2007 05:08:08 GMT

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

On Tue, Apr 03, 2007 at 09:04:59PM -0700, Paul Menage wrote:

> Have you posted the cpuset implementation over your system yet?

Yep, here:

<http://lists.linux-foundation.org/pipermail/containers/2007-March/001497.html>

For some reason, the above mail didnt make it into lkml (maybe it exceeded the max size allowed). I also have a updated version of that which I hope to post as soon as I am done with something else I am working on (sigh ..)

> The drawback to that is that every subsystem has to add a dentry to
> its state, and handle the processing.

Again this depends on whether every subsystem need to be able to support the user-space query you pointed out.

> >Do you see similar queries coming in for every resource controller object
> >(show me the path of cpu_acct, cpu_ctl, rss_ctl ... objects to which this
> >task belongs)? IMO that will not be the case, in which case we can avoid
> >adding N pointers (N = max hierarchies) in nsproxy just to support queries
> >of

> >those sort.

>

> OK, I see your argument that putting it in the aggregator probably
> isn't the best thing to do from a space point of view in the case when
> the number of aggregators

Sorry that sentence seems to be garbled by some mail router :)

Did you mean to say "when the number of aggregators sharing the same container object are more" ?

I agree ..Putting N pointers in container_group object just to support queries isn't justified at this point, because we don't know whether all subsystems need to support such queries.

> This seems like a place where my container_subsys_state object is
> useful - it can store a pointer to the container object (and be
> maintained by the generic container system), at a space cost of 1
> pointer per subsystem grouping, rather than N pointers per aggregator.

Yes that would be better than having N pointers in aggregator. From supporting purely user-space query pov, I think that is roughly same as having a 'dentry pointer' per resource object (what I mentioned earlier).

IMO we should add a dentry/container_subsys_state pointer only for those subsystems which need to support such queries ..

--

Regards,
vatsa
