
Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by [serue](#) on Mon, 12 Mar 2007 15:56:43 GMT

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

Quoting Srivatsa Vaddagiri (vatsa@in.ibm.com):

> On Fri, Mar 09, 2007 at 02:09:35PM -0800, Paul Menage wrote:

> > > 3. This next leads me to think that 'tasks' file in each directory doesnt make
> > > sense for containers. In fact it can lend itself to error situations (by
> > > administrator/script mistake) when some tasks of a container are in one
> > > resource class while others are in a different class.

> > >

> > > Instead, from a containers pov, it may be usefull to write
> > > a 'container id' (if such a thing exists) into the tasks file
> > > which will move all the tasks of the container into
> > > the new resource class. This is the same requirement we
> > > discussed long back of moving all threads of a process into new
> > > resource class.

> >

> > I think you need to give a more concrete example and use case of what
> > you're trying to propose here. I don't really see what advantage
> > you're getting.

>

> Ok, this is what I had in mind:

>

>

> mount -t container -o ns /dev/namespace

> mount -t container -o cpu /dev/cpu

>

> Lets we have the namespaces/resource-groups created as under:

>

> /dev/namespace

> |-- prof

> | |-- tasks <- (T1, T2)

> | |-- container_id <- 1 (doesnt exist today perhaps)

> |

> |-- student

> | |-- tasks <- (T3, T4)

> | |-- container_id <- 2 (doesnt exist today perhaps)

>

> /dev/cpu

> |-- prof

> | |-- tasks

> | |-- cpu_limit (40%)

> |

> |-- student

> | |-- tasks

> | |-- cpu_limit (20%)

> |
> |
>
>
> Is it possible to create the above structure in container patches?
> /me thinks so.
>
> If so, then accidentally someone can do this:
>
> echo T1 > /dev/cpu/prof/tasks
> echo T2 > /dev/cpu/student/tasks
>
> with the result that tasks of the same container are now in different
> resource classes.

What's wrong with that?

> That's why in case of containers I felt we shouldn't allow individual tasks
> to be cat'ed to tasks file.
>
> Or rather, it may be nice to say :
>
> echo "cid 2" > /dev/cpu/prof/tasks
>
> and have all tasks belonging to container id 2 move to the new resource
> group.

Adding that feature sounds fine, but don't go stopping me from putting
T1 into /dev/cpu/prof/tasks and T2 into /dev/cpu/student/tasks just
because you have your own notion of what each task is supposed to be.
Just because they're in the same namespaces doesn't mean they should get
the same resource allocations. If you want to add that kind of policy,
well, it should be policy - user-definable.

-serge

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
