Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by Herbert Poetzl on Fri, 09 Mar 2007 21:52:18 GMT
View Forum Message <> Reply to Message

On Fri, Mar 09, 2007 at 11:49:08PM +0530, Srivatsa Vaddagiri wrote:
> On Fri, Mar 09, 2007 at 01:53:57AM +0100, Herbert Poetzl wrote:
>>> The real trick is that I believe these groupings are designed to
>>> be something you can setup on login and then not be able to switch
>>> out of. Which means we can't use sessions and process groups as the
>>> grouping entities as those have different semantics.
>>
>> precisely, once you are inside a resource container, you
>> must not have the ability to modify its limits, and to
>> some degree, you should not know about the actual available
>> resources, but only about the artificial limits

the emphasis here is on 'from inside' which basically
boils down to the following:

 if you create a 'resource container' to limit the
 usage of a set of resources for the processes
 belonging to this container, it would be kind of
 defeating the purpose, if you'd allow the processes
 to manipulate their limits, no?

> From non-container workload management perspective, we do desire
> dynamic manipulation of limits associated with a group and also the
> ability to move tasks across resource-classes/groups.

the above doesn't mean that there aren't processes
outside the resource container which have the
necessary capabilities to manipulate the container
in any way (changing limits dynamically, moving
tasks in and out of the container, etc ...)

best,
Herbert

> --
> Regards,
> vatsa
> _____
> Containers mailing list
> Containers@lists.osdl.org
> https://lists.osdl.org/mailman/listinfo/containers

_____
Containers mailing list