
Subject: Re: [RFC] libcg: design and plans
Posted by [Dhaval Giani](#) on Wed, 05 Mar 2008 15:53:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, Mar 05, 2008 at 01:56:23PM +0200, Xpl++ wrote:

> Hi Dhaval,

>

> Dhaval Giani ???????:

>>> I was wonder if creating such library makes any sense at all, considering

>>> the nature of cgroups, the way they work and their possible application?

>>> It seems to me that any attempt to create a 'simple' API would actually

>>> result in something that will be much harder to use than just making raw

>>> mkdir/open/read/write/close operations.

>>>

>>

>> These simple APIs are nothing but raw mkdir/open/read/write/close

>> operations.

>>

>>

> Then why do you have to make the api? Everybody knows how to user these

> syscalls :)

> If its only this -> why use it? If it's not -> it aint simple anymore?

Well, that's just a small subset of the APIs. There are a lot of other things that libcg does.

>>> Another thing is suggested config for this lib would be more appropriate
>>> for a daemon rather than a library.

>>> In general - cgroup is a very flexible subsystem that can be used in a

>>> wide variety of ways and modes and trying to create a universal simple

>>> API would more likely result in something hard to manage and work with.

>>> The idea of having multiple container managers (applications that use

>>> libcg) creates a lots of questions and possible issues. Containers are

>>> supposed provide a flexible resource management and task grouping

>>> ability, which somewhat implies that there cannot be two different

>>> resource managers per node without one knowing well the

>>> desires/goals/methods of the other and vice versa. And should there be

>>> only one manager per node - probably it will be easier for it to use

>>> cgroup subsystem directly rather than using a wrapper library?

>>>

>>

>> I disagree. Allowing multiple resource managers allows more flexibility.

>> One thing the configuration subsystem aims to do is to allow permissions

>> to the groups. With this happening, a resource manager does not need to

>> about the existence of other groups, it can control only the resources

>> allotted to it. And since the top level is controlled by the

>> administrator, he is aware of the groups and their needs. (I've given an

>> example of such a scenario in the document as well).

>>
> Imagine having a shared/joint household savings account with your wife, and
> taking money from it without your wife knowing and vice versa .. then at
> some point when you thought that you have some \$5K to buy the new super
> duper laptop you dreamt about your entire life - surprise - no enough
> resources :)
> This is somewhat the equivalent of multiple independent resource managers
> :) It won't end well :)
> Should they be expected to be adequate in doing their job, they cannot be
> independent since they manage a shared resource.
>

I don't quite agree with your analogy here. The point here is that each resource manager operates in its own area, and has already been assigned some resources which it cannot change. Its more like you can take \$x at the most from the account and your wife \$y with $x+y \leq \text{total money}$.

--
regards,
Dhaval

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