
Subject: Re: Re: [RFC] libcg: design and plans
Posted by [Balbir Singh](#) on Wed, 05 Mar 2008 11:48:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

Denis V. Lunev wrote:

> On Tue, 2008-03-04 at 22:15 -0800, Paul Menage wrote:

>> Hi Dhaval,

>>

>> On Tue, Mar 4, 2008 at 7:23 AM, Dhaval Giani <dhaval@linux.vnet.ibm.com> wrote:

>>> Hi,

>>>

>>> We have been working on a library for control groups which would provide
>>> simple APIs for programmers to utilize from userspace and make use of
>>> control groups.

>>>

>>> We are still designing the library and the APIs. I've attached the
>>> design (as of now) to get some feedback from the community whether we
>>> are heading in the correct direction and what else should be addressed.

>> There are a few things that it would be nice to include in such a
>> library, if you're going to develop one:

>>

>> - the ability to create abstract groups of processes, and resource
>> groups, and have the ability to tie these together arbitrarily. E.g
>> you might create abstract groups A, B and C, and be able to say that A
>> and B share memory with each other but not with C, and all three
>> groups are isolated from each other for CPU. Then libcg would mount
>> different resource types in different cgroup hierarchies (you would
>> probably tell it ahead of time which combinations of sharing you would
>> want, in order that it could minimize the number of mounted
>> hierarchies). When you tell libcg to move a process into abstract
>> group A, it would move it into the appropriate resource group in each
>> hierarchy.

>

> There is one more important thing. In addition to the processes you must
> unite or provide a way to unite other objects like sockets. This is
> needed to create a group-based socket buffer management.

>

> The mapping between socket and a process does not exist right now and,
> we can have (virtually), sockets from different namespaces in one
> process.

>

Not sure how any of this is related to the library design we are discussing.

Your talking about writing a controller that groups based on sockets, that is a totally different thing.

--

Warm Regards,

Balbir Singh
Linux Technology Center
IBM, ISTL

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