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

Posted by [Paul Jackson](#) on Fri, 09 Mar 2007 19:36:23 GMT

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

Herbert wrote (and vatsa quoted):

- > 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

Not necessarily. Depending on the resource we are managing, and on how all encompassing one chooses to make the virtualization, this might be an overly Draconian permission model.

Certainly in cpusets, any task can see the the whole system view, if there are fairly typical permissions on some /dev/cpuset files. A task might even be able to change those limits for any task on the system, if it has stronger privilege.

Whether or not to really virtualize something, so that the contained task can't tell it is in side a cacoon, is a somewhat separate question from whether or not to limit a tasks use of a particular precious resource.

--

I won't rest till it's the best ...
Programmer, Linux Scalability
Paul Jackson <pj@sgi.com> 1.925.600.0401

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
