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

Posted by [Paul Jackson](#) on Mon, 12 Mar 2007 20:26:32 GMT

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

vatsa wrote:

> This assumes that you can see the global vfs namespace right?

>

> What if you are inside a container/vserver which restricts your vfs

> namespace? i.e /dev/cpusets seen from one container is not same as what

> is seen from another container .

Well, yes. But that restriction on the namespace is no doing of cpusets.

It's some vfs namespace restriction, which should be an orthogonal mechanism.

Well, it's probably not orthogonal at present. Cpusets might not yet handle a restricted vfs name space very well.

For example the /proc/<pid>/cpuset path, giving path below /dev/cpuset of task pid's cpuset, might not be restricted. And the set of all CPUs and Memory Nodes that are online, which is visible in various /proc files, and also visible in ones top cpuset, might be inconsistent if restricted vfs namespace mapped you to a different top cpuset.

There are probably other loose ends as well.

--

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>
