Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by Paul Menage on Wed, 07 Mar 2007 17:29:12 GMT
View Forum Message <> Reply to Message

On 3/7/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:
>
> > - when you do sys_unshare() or a clone that creates new namespaces,
> > then the task (or its child) will get a new nsproxy that has the rcfs
> > subsystem state associated with the old nsproxy, and one or more
> > namespace pointers cloned to point to new namespaces. So this means
> > that the nsproxy for the task is no longer the nsproxy associated with
> > any directory in rcfs. (So the task will disappear from any "tasks"
> > file in rcfs?)
>
> it "should" disappear yes, although I haven't carefully studied the
> unshare requirements yet.

That seems bad. With the current way you're doing it, if I mount hierarchies A and B on /mnt/A and /mnt/B, then initially all tasks are in /mnt/A/tasks and /mnt/B/tasks. If I then create /mnt/A/foo and move a process into it, that process disappears from /mnt/B/tasks, since its nsproxy no longer matches the nsproxy of B's root container. Or am I missing something?

Paul

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers