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

Posted by [serue](#) on Thu, 22 Mar 2007 14:39:09 GMT

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

Quoting Srivatsa Vaddagiri (vatsa@in.ibm.com):

> On Fri, Mar 09, 2007 at 10:50:17AM -0600, Serge E. Hallyn wrote:
> > The nsproxy container subsystem could be said to be that unification.
> > If we really wanted to I suppose we could now always mount the nsproxy
> > subsystem, get rid of tsk->nsproxy, and always get thta through it's
> > nsproxy subsystem container. But then that causes trouble with being
> > able to mount a hierarachy like
> >
> > mount -t container -o ns,cpuset
>
> What troubles will mounting both cpuset and ns in the same hierarchy
> cause?

Wow, don't recall the full context here. But at least with Paul's container patchset, a subsystem can only be mounted once. So if the nsproxy container subsystem is always mounted by itself, then you cannot remount it to bind it with cpusets.

> IMO that may be a good feature by itself, which makes it convenient to
> bind different containers to different cpusets.

Absolutely.

-serge

> In this case, we want 'ns' subsystem to override all decisions wrt
> mkdir of directories and also movement of tasks b/n different
> groups. This is automatically accomplished in the patches, by having ns
> subsystem veto mkdir/can_attach request which aren't allowed as per
> namespace semantics (but which may be allowed as per cpuset semantics).
>
> > so we'd have to fix something. It also slows things down...
>
> --
> Regards,
> vatsa

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

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

Posted by [Srivatsa Vaddagiri](#) on Thu, 22 Mar 2007 14:56:49 GMT

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

On Thu, Mar 22, 2007 at 09:39:09AM -0500, Serge E. Hallyn wrote:

> > What troubles will mounting both cpuset and ns in the same hierarchy
> > cause?

>

> Wow, don't recall the full context here.

Sorry to have come back so late on this :)

> But at least with Paul's container patchset, a subsystem can only be mounted
> once. So if the nsproxy container subsystem is always mounted by itself,
> then you cannot remount it to bind it with cpusets.

Sure ..I was thinking of mounting both ns and cpuset in same hierarchy
for first time itself.

--

Regards,
vatsa

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
