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

Posted by ebiederm on Thu, 08 Mar 2007 02:25:54 GMT

View Forum Message <> Reply to Message

"Paul Menage" <menage@google.com> writes:

> On 3/7/07, Eric W. Biederman <ebiederm@xmission.com> wrote:
>> The real trick is that I believe these groupings are designed to be something
>> you can setup on login and then not be able to switch out of.
>
> That's going to to be the case for most resource controllers - is that
> the case for namespaces? (e.g. can any task unshare say its mount
> namespace?)

With namespaces there are secondary issues with unsharing. Weird things like a simple unshare might allow you to replace /etc/shadow and thus mess up a suid root application.

Once people have worked through those secondary issues unsharing of namespaces is likely allowable (for someone without CAP_SYS_ADMIN). Although if you pick the truly hierarchical namespaces the pid namespace unsharing will simply give you a parent of the current namespace.

For resource controls I expect unsharing is likely to be like the pid namespace. You might allow it but if you do you are forced to be a child and possible there will be hierarchy depth restrictions. Assuming you can implement hierarchical accounting without to much expense.

Eric

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