
Subject: Re: [ckrm-tech] [PATCH 00/10] Containers(V10): Generic Process Containers

Posted by [serue](#) on Fri, 08 Jun 2007 18:08:37 GMT

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

Quoting Paul Menage (menage@google.com):

> On 6/8/07, Serge E. Hallyn <serge@hallyn.com> wrote:

> >

> >The problem is container_clone() doesn't call ->create explicitly, it

> >does vfs_mkdir. So we have no real way of passing in clone_task.

> >

>

> Good point.

>

> Looking at vfs_mkdir(), it's pretty simple, and really the only bits

> that apply to container_clone() are the call to ->mkdir() and possibly

> the call to fsnotify_mkdir(). (I think that's maybe how you did it

> originally?)

Yes it was.

> Maybe it would make sense to just call container_create() at that

> point directly, which would allow us more parameters.

I do fear that that could become a maintenance nightmare. For instance right now there's the call to fsnotify_mkdir(). Other such hooks might be placed at vfs_mkdir, which we'd then likely want to have placed in our container_mkdir() and container_clone() fns. And of course may_create() is static inline in fs/namei.c. It's trivial, but still if it changes we'd want to change the version in kernel/container.c as well.

What would be the main advantage of doing it this way? Do you consider the extra subys->auto_setup() hook to be avoidable bloat?

thanks,

-serge
