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

Posted by [Paul Menage](#) on Fri, 08 Jun 2007 18:13:41 GMT

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

On 6/8/07, Serge E. Hallyn <serue@us.ibm.com> wrote:

>
> 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.

Do we need to actually need to respect may_create() in
container_clone()? I guess it would provide a way for root to control
which processes could unshare namespaces.

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

I was thinking that it would be nice to be able to atomically set up
the resources in the new container at the point when it's created
rather than later. But I guess this way can work too. Can we call it
something like "clone()" rather than "auto_setup()"?

Paul
