
Subject: Re: [RFC][PATCH 1/5] Virtualization/containers: startup

Posted by [dev](#) on Mon, 06 Feb 2006 19:30:31 GMT

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

> As someone said to me a little bit ago, for migration or checkpointing
> ultimately you have to capture the entire user/kernel interface if
> things are going to work properly. Now if we add this facility to
> the kernel and it is a general purpose facility. It is only a matter
> of time before we need to deal with nested containers.

Fully virtualized container is not a matter of virtualized ID - it is
the easiest thing to do actually, but a whole global problem of other
resources virtualization. We can omit ID for now, if you like it more.

> Not considering the case of having nested containers now is just foolish.

> Maybe we don't have to implement it yet but not considering it is silly.

No one told that it is not considered. In fact PID virtualization send
both by IBM/us is abstract and doesn't care whether containers are
nested or not.

> As far as I can tell there is a very reasonable chance that when we
> are complete there is a very reasonable chance that software suspend
> will just be a special case of migration, done complete in user space.
> That is one of the more practical examples I can think of where this
> kind of functionality would be used.

Kirill
