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

Posted by [ebiederm](#) on Mon, 06 Feb 2006 18:37:07 GMT

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

Dave Hansen <haveblue@us.ibm.com> writes:

> On Mon, 2006-02-06 at 02:19 -0700, Eric W. Biederman wrote:
>> That you placed the namespaces in a separate structure from
>> task_struct.
>> That part seems completely unnecessary, that and the addition of a
>> global id in a completely new namespace that will be a pain to
>> virtualize
>> when it's time comes.
>
> Could you explain a bit why the container ID would need to be
> virtualized?

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.

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.

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.

Eric
