

---

Subject: Re: [PATCH] containers: define a namespace container subsystem

Posted by [serue](#) on Fri, 02 Feb 2007 17:23:08 GMT

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

---

Quoting Cedric Le Goater (clg@fr.ibm.com):

>  
> > The next steps are (not necessarily in order):  
> >  
> > 1. allow rm -rf to kill all processes under a  
> > ns\_container - with the intent of killing all  
> > processes in a virtual server  
> >  
> > 2. implement transitioning into a populated container,  
> > with the effect of setting the task's nsproxy to  
> > the one represented by the container.  
> >  
> > 3. define a file for each type of namespace in each  
>  
> could that file be a directory exposing some critical data  
> from each namespace ?

it probably could be, but that might be confusing since subcontainers are also directories. Would just putting the data into the namespace files suffice? This isn't sysfs so no 1-value-per-file restrictions...

> I would imagine the network devices for the net namespace  
> and be able to interact with them (Daniel ?). the task list  
> for the pid namespace, etc.

Well the tasklist will already be in the 'tasks' file created by the containers code :)

But actually, making them directories might actually be easier, because iirc cftypes only have f\_ops right now, whereas dirs already have i\_ops, so doing the symlink magic should be easier that way.

Great idea! :)

thanks,  
-serge

> > ns\_container, with the i\_op->symlink() defined to  
> > allow creation of a new ns\_container which references  
> > only some of the namespace pointers of an existing  
> > (child) container. All other namespaces will be  
> > taken from the existing process. In this way it  
> > is possible to enter just a network namespace of  
> > some vserver.

> > 4. probably make containers mac-aware, that is add a  
> > ->security pointer, and LSM hooks at appropriate  
> > points so that, for instance, SELinux can control  
> > vserver kill and enters.  
> >

---

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>

---