

---

Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem

Posted by [Srivatsa Vaddagiri](#) on Tue, 03 Apr 2007 14:09:24 GMT

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

---

On Mon, Apr 02, 2007 at 12:02:35PM -0600, Eric W. Biederman wrote:

> > If we loose directories, then we don't have a way to manage the  
> > task-group it represents thr' the filesystem interface, so I consider  
> > that bad. As we agree, this will not be an issue if initrd  
> > mounts the ns hierarchy atleast at bootup.

>

> I suspect that could be a problem if we have recursive containers.  
> Even by having a separate mount namespace for isolation you really  
> don't want to share the mount. If you can see all of the processes  
> you do want to be able to see and control everything.

Won't there be some master (VFS) namespace which can see everything? The idea would be then to list all containers in that namespace. I am visualizing that a master namespace listing all containers like that will be like a management console, from which you can monitor/control resource consumption of all containers.

I agree the individual containers themselves should not be able to mount and view other containers in this container/resource-control filesystem. I presume existing VFS namespace mechanism would enforce that restriction.

> I guess I want to ask before this gets to far. Why are all of the  
> namespaces lumped into one group?

I don't think they are. From Serge's patches, a new group (or a directory in container filesystem) is created everytime a new nsproxy is created (copy\_namespaces/sys\_unshare).

> I would think it would make much  
> more sense to treat each namespace individually (at least for the  
> user space interface).

--

Regards,  
vatsa

---