

---

Subject: Re: Re: [RFC][PATCH 0/2] user namespace [try #2]  
Posted by [Herbert Poetzl](#) on Thu, 07 Sep 2006 15:48:57 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Thu, Sep 07, 2006 at 07:40:23PM +0400, Kirill Korotaev wrote:

> > Here's a stab at semantics for how to handle file access. Should be  
> > pretty simple to implement, but i won't get a chance to implement this  
> > week.  
> >  
> > At mount, by default the vfsmount is tagged with a uid\_ns.  
> > A new -o uid\_ns=<pid> option instead tags the vfsmount with the uid\_ns  
> > belonging to pid <pid>. Since any process in a descendent pid  
> > namespace should still have a valid pid in the ancestor  
> > pidspace, this should work fine.  
> > At vfs\_permission, if current->nsproxy->uid\_ns != file->f\_vfsmnt->uid\_ns,  
> > 1. If file is owned by root, then read permission is granted  
> > 2. If file is owned by non-root, no permission is granted  
> > (regardless of process uid)  
> >  
> > Does this sound reasonable?

> imho this is acceptable for OpenVZ as it makes VE files to be  
> inaccessible from host. At least this is how I understand your  
> idea... Am I correct?  
>  
> > I assume the list of other things we'll need to consider includes  
> > signals between user namespaces  
> > keystore  
> > sys\_setpriority and the like  
> > I might argue that all of these should be sufficiently protected  
> > by proper setup by userspace. Can you explain why that is not  
> > the case?

> The same requirement (ability to send signals from host to VE)  
> is also applicable to signals.

at some point, we tried to move all cross context  
signalling (from the host to the guests) into a special  
context, but later on we moved away from that, because  
it was much simpler and more intuitive to handle the  
signalling with a separate syscall command

what I want to point out here is, that things like  
sending signals across namespaces is something which  
is not required to make this work

best,  
Herbert

> Thanks,  
> Kirill  
>  
> \_\_\_\_\_  
> Containers mailing list  
> Containers@lists.osdl.org  
> <https://lists.osdl.org/mailman/listinfo/containers>

---