Subject: Re: Re: [RFC][PATCH 0/2] user namespace [try #2]
Posted by Herbert Poetzl on Tue, 12 Sep 2006 14:07:08 GMT
View Forum Message <> Reply to Message

On Tue, Sep 12, 2006 at 05:52:40PM +0400, Kirill Korotaev wrote:
> Herbert Poetzl wrote:
> > On Thu, Sep 07, 2006 at 08:09:38PM +0400, Kirill Korotaev wrote:

> >>>>imho this in acceptable for OpenVZ as makes VE files to be
> >>>>inaccessiable 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

> >>I'm not sure what a separate context is for, but a separate syscall
> >>is definetely not a good idea.

> > care to explain _why_ you think so?
> cause duplicating syscalls with the same meaning but just working in a
> bit different situations doesn't look good.

hmm ... well, I guess the kernel doesn't look too good then :)

```
 .long sys_setuid16
 .long sys_getuid16
 .long sys_geteuid16
 .long sys_setreuid16 /* 70 */
 .long sys_setfsuid16
 .long sys_setresuid16
 .long sys_getresuid16 /* 165 */
 .long sys_getuid
 .long sys_geteuid
 .long sys_setreuid
 .long sys_setresuid
```

```
.long sys_getresuid
.long sys_setuid
.long sys_setfsuid /* 215 */

.long sys_umount /* recycled never used phys() */
.long sys_oldumount

.long sys_olduname
.long sys_uname
.long sys_newuname

.long sys_old_getrlimit
.long sys_getrlimit
```

best,
Herbert

> Kirill
