
Subject: Re: unshare() pid ns
Posted by [serue](#) on Thu, 30 Aug 2007 14:35:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Pavel Emelyanov (xemul@openvz.org):
> Serge E. Hallyn wrote:
> > Quoting Pavel Emelyanov (xemul@openvz.org):
> > > sukadev@us.ibm.com wrote:
> > > > Pavel,
> > > > unshare() of pid ns seems to fail with -EINVAL in 2.6.23-rc3-mm1.
> > > > I thought we supported it in the earlier patchsets. I guess
> > > > I missed that in the review of recent patchsets.
> > > I disabled unsharing of pid namespaces because it's almost
> > > impossible. Look - you have to reattach all the pids to the
> > > task with saving its ids as seen in previous namespaces.
> >
> > We agree, but thought you for some perverse reason preferred unshare to
> > clone for pidns :)
>
> I did that in my first version of patches, but then realized
> that such problem (the need in reattaching pids) makes the
> unsharing ugly.
>
> BTW, unsharing of a pid namespace is a valid operation, so I
> think I will enable it in the nearest future. I have some
> thought on how to make such a reattach ;)

Alrighty :)

Of course it's not just the kernel ugliness, but also the userspace ugliness, for instance the rumored (I haven't looked to confirm) caching of pids by glibc.

But in the end if we can achieve symmetry between all the CLONE_NEW* flags all the better.

-serge

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
