
Subject: Re: [PATCH 0/4] Helper patches for PTY namespaces

Posted by [serue](#) on Sun, 13 Apr 2008 00:59:36 GMT

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

Quoting Eric W. Biederman (ebiederm@xmission.com):

> On Sat, 2008-04-12 at 10:29 -0700, sukadev@us.ibm.com wrote:

> > Some simple helper patches to enable implementation of multiple PTY

> > (or device) namespaces.

> >

> > [PATCH 1/4]: Propagate error code from devpts_pty_new

> > [PATCH 2/4]: Factor out PTY index allocation

> > [PATCH 3/4]: Move devpts globals into init_pts_ns

> > [PATCH 3/4]: Enable multiple mounts of /dev/pts

> >

> > This patchset is based on earlier versions developed by Serge Hallyn

> > and Matt Helsley.

>

> Suka Stop.

>

> The first two patches appear to just be cleanups and as such should be

> able to stand on their own. Mentioning that you found these

> opportunities while working on your pts namespace is fine. Justifying

> the cleanups this way is not.

>

> When you mentioned you intended to just resend the cleanups this is not

> at all what I thought you were about. So please just get the first two

> patches in a form that stands by themselves.

It took me a minute to figure out what you were offended by, but I see, the patches still introduce a "pts ns", which shouldn't exist at all.

> The pts namespace as designed is not acceptable.

>

> The problem you are trying to solve with the pts namespace is real.

>

> So what we need is a device namespace and possibly and incremental path

> to get there. A device namespace would abstract the device number to

> device mapping for all devices. A safe incremental path would disable

> all device number to device mappings for process in that namespace. So

> all functionality would be gone, then it would enable certain mappings

> and certain pieces of the functionality if the code in the kernel is not

> clean enough that we can do it all in one go.

>

> We need to see that path. Only then can we take patches that add

> namespace specific goo. The pattern I am proposing has worked quite

> well for the network namespace. Meanwhile the uid namespace which

> follows the pattern you seem to be following now seems does not look to

> be completed any time soon.

(you know perfectly well that we're holding off on any further real uidns work until netns is complete and you have time to partake in the design - it wouldn't be any fun without you :)

> So send your clean up patches and then let's architect this thing so we
> are really talking about a namespace. Then we can update devpts to
> capture the device namespace on mount and do it's work in a namespace
> specific manner.

Sounds reasonable to me.

thanks,
-serge

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