
Subject: Re: [RFC][PATCH] Do not set /proc inode->pid for non-pid-related inodes
Posted by [serue](#) on Wed, 21 Mar 2007 01:02:49 GMT

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

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

> Dave Hansen <hansendc@us.ibm.com> writes:

>

> > On Tue, 2007-03-20 at 09:51 -0600, Eric W. Biederman wrote:

> >> Outlive is the wrong concept. Ideally we want something that will

> >> live as long as there are processes in the pid_ns.

> >

> > How about they just live as long as there is a mount? Now that we _can_

> > have multiple superblocks and meaningful vfsmounts, I think it's time to

> > make it act like a normal filesystem.

>

> Agreed.

>

> My concern is that the mount will outlive the pid namespace. In which

> case we need something that is safe to test when the pid namespace goes

> away.

Offhand I would assume the mount would get a reference to the pidns.

pidns may be empty, but would exist.

-serge

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
