Subject: Re: [RFC][PATCH] Do not set /proc inode->pid for non-pid-related inodes Posted by serue on Mon, 26 Mar 2007 18:36:13 GMT

View Forum Message <> Reply to Message

```
Quoting Eric W. Biederman (ebiederm@xmission.com):
> Dave Hansen <hansendc@us.ibm.com> writes:
> > On Mon, 2007-03-26 at 11:12 -0600, Eric W. Biederman wrote:
> >>
>>> In (at least one version of) Dave's patches, the /proc your pidns is
>>> automatically used when you use /proc. In that case a /proc should
>>> just go away when the last task goes away, since noone else can use
> >> > that /proc.
> >>
>>> Unless I am rather confused that does extremely nasty things to
>>> the VFS dentry cache. Because a dentry can point at one process
>>> one minute and another process the next. It is doable but only
>>> at the cost of decreased performance.
> >
>> By using, I think Serge meant "mounting". We're going to statically
>> assign a /proc mount to a namespace when the mount is created, not fudge
> > it around at runtime.
> I think Serge misread that. He specifically talked about not having
> to remount /proc. Therefore I believe he understood it did something
> at runtime.
Yes, looks like I misread the code.
I thought a new sb was created automatically if needed. Going back I
see this is only done at mount time.
> > How does this thrash the dcache?
> Fudging at runtime is horrible.
> Caching the pid_mnt in the pid_namespace appears fine, as long as
> we get the reference counting correct.
> Eric
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers
```