
Subject: Re: [RFC][PATCH] Do not set /proc inode->pid for non-pid-related inodes
Posted by [serue](#) on Mon, 26 Mar 2007 14:13:48 GMT

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

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

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

>

> > So, doesn't that problem go away (or at least move to be umount's duty)

> > if we completely disconnect those inodes' lifetime from that of any

> > process or pid namespace?

>

> If the last process has exited the pid namespace I would like the

> code to continue to behave as it currently does.

>

> I would like readdir on /proc/ to not even try to show any pids when

> there are no pids or pid related files in the pid namespace.

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.

I like that behavior, because otherwise (a) we require every new pid_namespace to start by remounting /proc ere they get undefined behavior, and (b) to gain anything from it, we would need a way to refer to another pidspace for the sake of mounting it's proc, i.e.

```
mount -t proc -o init_pid=7501 proc_vserver1 /vserver1/proc
```

-serge

> I would like /proc/self to completely disappear when there are not

> any pids in the pid namespace.

>

> I misspoke in when I said that /proc/<pid> was affected. The function

> is proc_pid_readdir and it is a subset of /proc/ so it gets a little

> confusing.

Containers mailing list

Containers@lists.linux-foundation.org

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