Subject: Re: [RFC][PATCH 03/14] use pid_nr in procfs
Posted by serue on Wed, 21 Mar 2007 17:09:32 GMT
View Forum Message <> Reply to Message

Quoting Eric W. Biederman (ebiederm@xmission.com):
> sukadev@us.ibm.com writes:
>
> > From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
> > Subject: [RFC][PATCH 03/14] use pid_nr in procfs
> >
> > With containers, a process can have different pid_t values in different
> > pid namespaces. To ensure we get the correct pid_t value in any context,
> > we should use pid_nr() function rather than directly accessing either
> > task->pid or pid->nr.
>
> To clarify my previous comment. I believe to get this right
> we need a factor of pid_nr:
> pid_t __pid_nr(struct pid_namespace *ns, struct pid *pid)
>
> That we can pass the pid_namespace from the super block of the
> mount to handle the proc case.
>
> This is a case I have wanted to avoid but in this instance I don't
> see any other way to get the code correct.
>
> i.e. Since the pid namespace of /proc gets set at mount time who
> you are should not vary the result it gives.

Seems to make sense.

However since meaningful /proc usage under multiple pid namespaces will
require Dave's patch to accomodate multiple proc_mnt's, and since that
patch seems to keep having to change as this pid_ns patchset changes,
how about we drop this patch for now and have (a corrected version of)
it included in Dave's /proc patchset on top of the rest of this set?

-serge

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