Subject: Re: Q: How complete is the pid namespace in mainline Posted by ebiederm on Fri, 26 Oct 2007 09:33:25 GMT

View Forum Message <> Reply to Message

Cedric Le Goater <clg@fr.ibm.com> writes:

- > Eric W. Biederman wrote:
- >> Guys how complete do you fee the pid namespace support is that
- >> has been merged into Linus's tree?

>>

- >> My impression until I started reading through code earlier today
- >> was that the support was just about done except for a couple of
- >> tricky details.

>

> Yes It looks sane.

>

> Here's what I have in mind:

>

- > * there are some patches from suka that make sure the pid namespace init
- > is not getting abusively killed by one of this children
- > * the pid cleanup is not complete
- > . locks
- > . kthread (i should work soon on improving kthread to support
- > signals)

>

- > IMO, the proc mount shouldn't be under the pid namespace. we will
- > need that sooner or later.

I was hoping to get a larger list of unfixed issues.

Currently from my review I have generated about 25 bugfix patches. Several of them in some fairly obvious places.

I think it is a good base to build on, but it feels to like we still have a significant ways to go.

I think the assumption that we can use global pid numbers instead of instead of struct pids is racy, and a serious maintenance problem. It leads to silent breakage of routines like get_net_ns_by_pid, and possibly a couple of other places.

I'm really not happy with pid_nr meaning a pid number in the init_pid_ns and pid_vnr meaning a pid number meaning a pid in the local pid namespace. But that is just a matter of names so I don't think it has caused any problems. Short of making it to easy to get a pid number in the &init_pid_ns.

Eric

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