Subject: Re: Q: How complete is the pid namespace in mainline Posted by ebiederm on Fri, 26 Oct 2007 23:09:26 GMT

View Forum Message <> Reply to Message

sukadev@us.ibm.com writes:

- > Eric W. Biederman [ebiederm@xmission.com] wrote:
- > | sukadev@us.ibm.com writes:
- > | >
- > | > Dave had suggested we print a warning the first time a container-init
- > forks()
- > | > without a handler for a fatal signal. I was planning on adding that as
- > | > patch 4 of the signal patch set and get some feedback.

> |

- > | Yes. How to cleanly handle signalling of container init is
- > | a tricky one. It does sound like you have made a reasonable start
- > | there.

> |

- > | Suka it is a lot more then that. How much more I'm not certain
- > | of. I suspect the only way to find the rest of the cases is
- > | just go through the code with a fine tooth come and read and look.

- > I agree. I did not mean to ignore the kthread conversions and was only
- > referring to the core pid namespace clone stuff.

Sure, and that make sense.

- > | So far doing that it has not at all hard for me to find either
- > | bugs or places where the implementation can be improved.

>

- > | Currently we have little things like kill(-1,...) signalling the
- > | wrong set of processes, and a couple of proc bugs.

- > I just realized the fix for this is in the signal patchset I was
- > referring to.

> https://lists.linux-foundation.org/pipermail/containers/2007-August/006987.html

> I notice that you have sent a patch for the kill -1.

Yes. I'm trying to get out as many simple little bug fixes as I can.

Sorry for missing the fact you guys had generated some patches to address this. Still I think mine is a little more comprehensive and shorter;)

That bug is on my list of really nasty bugs I want to avoid.

- > The proc_mnt bug Linus found seems to have slipped through when
- > merging Pavel's and my patches.

I really don't mind a handful of little bugs, it would be surprising if something hadn't slipped through at this point.

As long as everyone is aware that it is going to take a bit to find everything and stabilizing it all and everyone keeps looking we should be fine.

Oh. Do you know if there was a good reason for forcing the tty, session, and process group of a the first process in a pid namespace?

Eric

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