Subject: Re: Q: How complete is the pid namespace in mainline Posted by Sukadev Bhattiprolu on Fri, 26 Oct 2007 17:17:18 GMT

View Forum Message <> Reply to Message

Eric W. Biederman [ebiederm@xmission.com] 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.

The only thing that I know is pending is the issue of signalling container-init. We have not been able to find a clean fix for it.

The problem now is that a process in a child namespace can terminate its container-init and thereby the entire container. We have a 3-patch set (Oleg's and mine) that kind of addresses this. The scenario where the patchset fails is:

- the container-init has a blockable, fatal signal blocked
- a descendant of the container-init posts the fatal signal to container-init.
- container-init then unblocks the signal without ignoring or handling the signal.

In this case again the container-init can be terminated.

(by fatal I mean a signal whose default action is to terminate the process SIGKILL is of couse not blockable and is not a problem)

This issue can be addressed in user-space by the container-init - which should just ignore the fatal signal or setup a handler for it.

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.

Suka

Contain are realling list

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