
Subject: Re: [RFC][PATCH] Do not set /proc inode->pid for non-pid-related inodes
Posted by [ebiederm](#) on Thu, 22 Mar 2007 15:10:28 GMT

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

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

>> Back to the main subject I still don't understand the idea of running
>> a kernel daemon as pid == 1. What would that buy us?
>
> mostly a child reaper when there are no /sbin/init but its pid cannot
> be 1.

Yes we should be able to assign just about any process as the child_reaper
for the pid namespace. That is an entirely distinct concept in the kernel.
Although we have to be careful we don't reap ourselves, so there are some
subtle cases but it isn't a pid issue.

If this is just an application container it might make sense to have a
kernel thread standing in for /sbin/init.

Basically the whole application container thing is something we need to
revisit when we have the basic pid namespace implemented and before we
finalize things, and remove the CONFIG_EXPERIMENTAL tag.

Eric

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