
Subject: Re: [RFC][PATCH] Do not set /proc inode->pid for non-pid-related inodes
Posted by [ebiederm](#) on Wed, 21 Mar 2007 16:57:24 GMT

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

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

> Eric W. Biederman wrote:

>> "Serge E. Hallyn" <serue@us.ibm.com> writes:

>>

>>> So how do you see us enforcing pid1's existence? Somehow keep it from

>>> fully exiting, or just kill all the processes in it's namespace if it

>>> exits?

>

> what about a kthread that would be spawned when a task is cloned in an

> unshared pid namespace ? This is an extra cost in term of tasks.

If you use kernel_thread this can happen. (Die kernel_thread).

If you use the kthread interface keventd will be the parent process and we won't have problems. Thus most users of kernel_thread need to be fixed to use the kthread interface.

Thanks for the reminder of this one, I had forgotten that bit of reasoning for updating kernel_thread users.

>> Killing all other processes in the namespace when pid1 exits is what

>> I implemented last time around.

>

> this looks like a sane thing to do.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
