
Subject: Re: [RFC][PATCH] Do not set /proc inode->pid for non-pid-related inodes
Posted by [ebiederm](#) on Mon, 26 Mar 2007 17:24:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

> The way I'm testing pidspaces right now is
>
> ns_exec -c -p /usr/sbin/sshd -p 9999
>
> in which case sshd is pid1. Works fine...
>
> Would it be very limiting to have the first process have to stick
> around? (I'm asking, not criticizing - it's *my* preference that we
> allow pid==1 to exit, but if that's really not advantageous then
> maybe it's not worth fixing the ugly pieces that require it right
> now - afaik right now that's only the fact that PROC_INODE(/proc)->pid
> points to the struct pid for pidnr==1)

The /proc part is easy to fix. All I want to see there is that we do the right thing with pid related files. When the pid namespace is empty.

The practical reason for only allowing a pid namespace while pid == 1 exists, is something much more simple.

pid == 1 must exist today. We get into an extension of the semantics if we allow the case where pid == 1 exists. Semantic extensions can be very tricky, and we are way too early to see what the impact of such a semantic extension would be.

Therefore I request that we get a correct and work pid namespace before we try and extend things.

I also request that until questions like this are settled we leave the whole thing CONFIG_EXPERIMENTAL.

I have yet to see how we are going to implement things such as kill -1. And the other changes. There are huge chunks of functionality that we haven't gotten to yet.

Eric

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