
Subject: Re: Processes with multiple pid_t values
Posted by [ebiederm](#) on Tue, 09 Jan 2007 04:22:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sukadev Bhattiprolu <sukadev@us.ibm.com> writes:

> Hmm. I have this mail sitting in my sent folder since Dec 18 but do
> not see it in the Containers archives. If you already received, sorry
> to spam. Appreciate any clarifications of my doubt :-)

There is something about your usage of the word doubt that does not quite feel correct. I wonder if it is just me.....

> For my example above, I guess you are saying we want the following ids.

>
> PID PPID PGID SID
>
> init pid ns 1234 1233 1230 1220
>
> child pid ns 1 0 1 1

>
>
> If so, my confusion is that for the process 1234, following are true
> right ?

>
> task_pgrp(current) != task_session(current) != task_pid(current)

>
> If we associate a list of [namespace, id] tuples with each struct pid,
> then three different struct pids would have the same id in the child
> namespace (unless the process calling clone() is both a session and
> process group leader).

I guess it depends on how we set it up, it all depends on how the groups are formed.

There is no reason to believe any particular relationship between the pid pgrp and session of the calling process from which we inherit these things. If we can make it all work without doing the equivalent of setsid() at clone time I think it is a better implementation.

However by default /sbin/init does call setsid() at which point we know we have the relationship: pid == pgrp == session.

So once setsid() is called (the common case) we will have the id relationship above. In fact we force that relationship for /sbin/init today at kernel startup time.

Eric

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
