
Subject: Re: [PATCH] allow a task to join a pid namespace
Posted by [Glauber Costa](#) on Tue, 05 Jun 2012 11:33:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 06/05/2012 01:36 PM, Daniel Lezcano wrote:
> On 06/04/2012 03:33 PM, Glauber Costa wrote:
>> Currently, it is possible for a process to join existing
>> net, uts and ipc namespaces. This patch allows a process to join an
>> existing pid namespace as well.
>>
>> For that to remain sane, some restrictions are made in the calling process:
>>
>> * It needs to be in the parent namespace of the namespace it wants to jump to
>> * It needs to sit in its own session and group as a leader.
>>
>> The rationale for that, is that people want to trigger actions in a Container
>> from the outside. For instance, mainstream linux recently gained the ability
>> to safely reboot a container. It would be desirable, however, that this
>> action is triggered from an admin in the outside world, very much like a
>> power switch in a physical box.
>>
>> This would also allow us to connect a console to the container, provide a
>> repair mode for setups without networking (or with a broken one), etc.
>
> Hi Glauber,
>
> I am in favor of this patch but I think the pidns support won't be
> complete and some corner-cases are not handled.
>
> May be you can look at Eric's patchset [1] where, IMO, everything is
> taken into account. Some of the patches may be already upstream.
>
> Thanks
> -- Daniel
>

Daniel,

Please let me know what you think of the attached patch. It is on top of Eric's tree that you pointed me to, but I am not really using any of its functionality, so this would be equally doable in current mainline kernel - but I wanted to make sure it integrates well with what Eric is doing as well.

It is a bit wasteful space-wise, but this approach pretty much guarantees we don't need to update pointers anywhere - therefore, totally lock-free. pid->level is only updated after switch_task_namespaces(), so every call before that will see correct

information up to the previous level.

File Attachments

1)

[0001-pid_ns-expand-the-current-pid-namespace-to-a-new-nam.patch](#), downloaded 486 times
