Subject: Re: [patch -mm 1/5] mqueue namespace : add struct mq_namespace
Posted by Cedric Le Goater on Wed, 03 Oct 2007 07:38:58 GMT
View Forum Message <> Reply to Message

[
  I have big fingers this morning and I managed to send this email
  while typing it ... see below for the end. I should be awake now :)
]

> The really challenging case to handle here is what happens if we are
> signaling to someone in a sibling pid namespace. What do we set the
> parent pid in the siginfo struct to. I think we agreed that 0 (blame
> the kernel) is the appropriate pid last time we talked about this.

0 seems appropriate for a signal coming from a parent namespace, yes. but
here, we could be sending a signal from a child or sibling namespace.

> I'm worried now that the concept of vpid has confused someone. It
> still doesn't feel right to me to call one pid value more or less
> virtual then any other so the concept of a virtual pid doesn't make
> sense to me. The way I have always thought of it is:
> - pid_nr(struct pid *)
>   The pid in the current pid namespace.
> - __pid_nr(struct pid_namespace, struct pid *)
>   The pid in some specified pid namespace.
>
> With struct pid being defined to be global and doing something
> appropriate in all pid namespaces.
>
> Thinking about this concern that Cedric raises is actually independent
> of the mqueue namespace and seems to be totally a pid namespace thing.
> Because the only way this happens if we happen to share the mqueue
> namespace. (i.e. what we are doing now).

Is there a way to catch this general issue (we have the same in sigio)
in the kill*(struct pid) routines ? spit a big warning when the
current->nsproxy->pid_ns is not a parent ?

Thanks !

C.

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