
Subject: Re: pspace child_reaper
Posted by [ebiederm](#) on Tue, 29 Aug 2006 15:24:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

> Hello All,
>
> Eric, in your initial proof of concept on the pid namespace, you were
> defining a child_reaper per pid namespace.
>
> IMO, we can't use init_task as a child_reaper in a pid namespace because we
> will have pid collision which might result in a breakage of the init_task.

The kernel doesn't use init_task (The idle thread) once it starts
init. Reaping children is the job of pid == 1.

> Here are some questions on the model you intended to follow :
>
> Do you think we should have a child_reaper task per container ?
We have an init per container so yes.

> Could we use a kthread to do the job ?
Definitely not.

> Could that kthread be global to all pid namespace ?
Makes no sense.

> Any completely different idea on the topic ?
Init reaps the children, and I believe there are parts of user space
that depend on this. We shouldn't change that semantic.

Eric

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