
Subject: pspace child_reaper

Posted by [Cedric Le Goater](#) on Tue, 29 Aug 2006 14:38:23 GMT

[View Forum Message](#) <> [Reply to Message](#)

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.

Here are some questions on the model you intended to follow :

Do you think we should have a child_reaper task per container ?

Could we use a kthread to do the job ?

Could that kthread be global to all pid namespace ?

Any completely different idea on the topic ?

thanks,

C.

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>

Subject: Re: pspace child_reaper

Posted by [ebiederm](#) on Wed, 30 Aug 2006 15:17:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

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

> I think there is some confusion here. we need a process 1 of course but it
> does not have to be necessarily a user space /sbin/init process.

>

>> I also believe that most of the advantages of not having an init
>> process can be had with a trivial (probably static) init program that
>> only calls waitpid. Taking essentially no resources.

>

> yes that could be one solution, or doing it in a kthread.

>

> but leaving it up to user space to do the right thing (handle sigchld) in
> the process doing the unshare seems a better solution.

I believe we are in violent agreement :)

Eric

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
