
Subject: Re: pspace child_reaper

Posted by [Herbert Poetzl](#) on Wed, 30 Aug 2006 17:11:39 GMT

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

On Wed, Aug 30, 2006 at 03:03:22PM +0200, Cedric Le Goater wrote:

> Serge E. Hallyn wrote:

> >>> But don't confuse /sbin/init with the reaper. /sbin/init will only
> >>> know about pids in it's own container, so if we do need to call
> >>> to userspace for every task which the reaper hears about, then
> >>> we will need a per-container reaper. But if that's not the case
> >>> (and i haven't seen where it is), then userspace is simply not
> >>> involved, and so one global reaper suffices, since it will know
> >>> not about pid_ts but about struct pids == (container,pid_t).

> >> First I don't remember the details but there was at least one
> >> case where the vserver guys hit an application that cared.

> > If that's the case, then per-container reaper it is.

first, to comment on that, here is what we (Linux-VServer)
do regarding init and reaper ...

the 'host' reaper is the original init process started
on the host. besides this one, we also know a guest
reaper, which 'usually' is identical with the guest's
init process, but, can be set to something differently.
most important, if that reaper is not present or defined
the kernel logic falls back to reparenting 'to-be-reaped'
children to the host init, which successfully reaps them

> > Herbert, do you recall which software that was?

no, I do not remember, but I will have a look in the
IRC logs, probably something will turn up

> I'm not sure herbert is on the list.

indeed, I am not (currently I'm in the progress of
subscribing to that mailing list too :)

best,
Herbert

> C.

Containers mailing list

Containers@lists.osdl.org

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