

---

Subject: Re: [RFC][PATCH] Add child reaper to struct pspace

Posted by [dev](#) on Fri, 08 Sep 2006 17:32:27 GMT

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

---

Sukadev Bhattiprolu wrote:

> Cedric Le Goater [[clg@fr.ibm.com](mailto:clg@fr.ibm.com)] wrote:

```
> |  
> | <snip>  
> |  
> | > */  
> | > static void  
> | > forget_original_parent(struct task_struct *father, struct list_head *to_release)  
> | > @@ -669,7 +670,7 @@ forget_original_parent(struct task_struct  
> | > do {  
> | >     reaper = next_thread(reaper);  
> | >     if (reaper == father) {  
> | > - reaper = child_reaper;  
> | > + reaper = father->pspace->child_reaper;  
> | >     break;  
> | > }  
> | > } while (reaper->exit_state);  
> | > @@ -857,7 +858,7 @@ fastcall NORET_TYPE void do_exit(long co  
> |  
> | what about killing all the task in that pid space if child_reaper == init  
> | dies ?  
> |  
> |  
> |
```

> We probably need that for instance when a process in the parent pspace  
> kills the init of a child pspace, we should destroy the child pspace  
> by killing all the tasks in the child pspace including the child reaper.  
exactly. the situation you described is how we do handle it.  
you can check `do_initproc_exit()` function in OpenVZ  
to check how it can be done and probably save some of your time.  
(<http://git.openvz.org/?p=linux-2.6-openvz;a=summary>)

> I guess we need to maintain a list of task\_structs in the pspace and walk  
> that list. Will work on that as a separate patch.  
wait. we either need to have a list of `_pids_` or it  
should be called `task_namespace`, not `pid`, since we are adding more  
code related to tasks.

Kirill

---