Subject: Re: [PATCH 3/5] Use task_pid() to find leader's pid Posted by Oleg Nesterov on Sun, 15 Jul 2007 12:17:50 GMT

View Forum Message <> Reply to Message

Sukadev Bhattiprolu wrote:

> Use task_pid() to get leader's pid since find_pid() cannot be used
> after detach_pid(). See comments in the code below for more details.
> ...
> + * Note: With multiple pid namespaces, active pid namespace of
> + * a process is stored in its struct pid. The detach_pid
> + * below frees the struct pid, so we will have no notion
> + * of an active pid namespace until we complete the
> + * subsequent attach_pid(). Which means - calls like
> + * find_pid()/pid_to_nr() return NULL and cannot be used
> + * between the detach_pid() and attach_pid() calls.

I think both the changelog and the comment are confusing,

```
> detach_pid(tsk, PIDTYPE_PID);
> tsk->pid = leader->pid;
> attach_pid(tsk, PIDTYPE_PID, find_pid(tsk->pid));
> + attach_pid(tsk, PIDTYPE_PID, task_pid(leader));
```

because the change itself looks like an obvious performance fix, even we don't use multiple pid namespaces. I don't think it is good idea to add a fat comment which doesn't match the current reality, and find_pid() should be avoided anyway.

Stupid question: why do we need to put the pid namespace into the struct pid? Isn't it better if the user of the struct pid should know its ns? For example, if /proc does put_pid(), that pid should be from the active namespace.

Sukadev, could you come if you do that kind of changes?

Oleg.

Containers mailing list

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