
Subject: Re: [RFC][PATCH 06/16] Define is_global_init()
Posted by [Pavel Emelianov](#) on Fri, 25 May 2007 06:39:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dave Hansen wrote:

```
> On Thu, 2007-05-24 at 13:24 +0400, Pavel Emelianov wrote:
>>> | > +int is_global_init(struct task_struct *tsk)
>>> | > +{
>>> | > + return (task_active_pid_ns(tsk) == &init_pid_ns && tsk->pid == 1);
>>> |
>>> | This can OOPS if you pass arbitrary task to this call...
>>> | tsk->nsproxy can already be NULL.
>>>
>>> Hmm. You are right. btw, this could be a bisect issue. Patch 9 of uses
>>> pid_ns from pid->upid_list and removes nsproxy->pid_ns.
>> Yes, but that patch is not good either.
>> task_pid(tsk) may become NULL as well and this will oops.
>
> Have you reviewed the call paths to make sure this can actually happen
> in practice?
>
> This just seems like another one of those racing-with-task-exit races.
> Shouldn't be too invasive to solve.
```

It is, but if we make patch that OOPSes the kernel in 0.1%
of cases and we do know this - this MUST be fixed.

```
> -- Dave
>
>
```

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
