
Subject: Re: [PATCH 2/3] Pid ns helpers for signals
Posted by [Oleg Nesterov](#) on Sat, 01 Sep 2007 11:56:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 09/01, Oleg Nesterov wrote:

```
>
> On 08/31, sukadev@us.ibm.com wrote:
> >
> > +static struct pid_namespace *get_task_pid_ns(struct task_struct *tsk)
> > +{
> > + struct pid *pid;
> > + struct pid_namespace *ns;
> > +
> > + pid = get_task_pid(tsk, PIDTYPE_PID);
> > + ns = get_pid_ns(pid_active_ns(pid));
> > + put_pid(pid);
> > +
> > + return ns;
> > +}
>
> Hmm. Firstly, we don't need this for the "current", but all users of this func
> also do get_task_pid_ns(current).
>
> Also, we don't need get/put_pid. rcu locks are enough,
>
> rcu_read_lock();
> ns = get_pid_ns(pid_active_ns(task_pid(tsk)));
> rcu_read_unlock();
>
> However, do we really need this complications right now? Currently, we use
> this "compare namespaces" helpers only when we know that "struct pid" is
> stable. We are sending the signal to that task, it must be pid_alive(), and
> we either locked the task itself, or we hold tasklist.
```

(forgot to mention)

Otherwise, it is not safe to use "tsk" in get_task_pid_ns(), so I don't think these get/put pid/pid_ns games make too much sense.

Oleg.

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