
Subject: Re: [PATCH 2/3] Pid ns helpers for signals
Posted by [Oleg Nesterov](#) on Sat, 01 Sep 2007 11:29:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 08/31, sukadev@us.ibm.com wrote:

```
>
> Define some helper functions that will be used to implement signal semantics
> with multiple pid namespaces.
>
> is_current_in_ancestor_pid_ns(task)
>
> TRUE iff active pid namespace of 'current' is an ancestor of
> active pid namespace of @task.
>
> is_current_in_same_or_ancestor_pid_ns(task)
>
> TRUE iff active pid namespace of 'current' is either same as
> or an ancestor of active pid namespace of @task.
```

These names are awful :) Yes, yes, it was me who suggested them... No, I can't suggest something better.

```
> + * Caller must hold a reference to @pid.
> + */
> +static inline struct pid_namespace *pid_active_ns(struct pid *pid)
> +{
> + if (!pid)
> + return NULL;
> +
> + return pid->numbers[pid->level].ns;
> +}
```

Well, the comment is a bit misleading. Yes, my previous comment was not very clear. Yes, the function itself is not safe unless you know what are you doing, like, for example, `get_pid()`. I think it is better to just kill the comment. Please see below.

```
> +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(tks)));  
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.

Oleg.

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