
Subject: Re: [PATCH 1/3] Signal semantics for /sbin/init
Posted by [Oleg Nesterov](#) on Mon, 03 Sep 2007 16:45:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 09/03, sukadev@us.ibm.com wrote:

```
>
> Oleg Nesterov [oleg@tv-sign.ru] wrote:
> | On 08/31, sukadev@us.ibm.com wrote:
> | >
> | > -static int sig_ignored(struct task_struct *t, int sig)
> | > + // Currently this check is a bit racy with exec(),
> | > + // we can _simplify_ de_thread and close the race.
> | > + if (likely(!is_container_init(tsk->group_leader)))
> | > + return 0;
> | > +
> | > + if (!in_interrupt())
> | > + return 0;
> |
> | I don't understand why you are trying to mix this patch with pid_ns changes.
> |
> | We don't need in_interrupt() check unless we use current do decide if the
> | signal goes from the parent namespace.
> |
> | And in fact, I'd personally prefer to use "is_global_init()" for this patch,
> | because it hopefully can fix the problems we have even without namespaces.
> | This also matches the current check in get_signal_to_deliver().
> |
> Sorry. I wasn't paying enough attention to this patch and including it
> only for reference. Was planning to replace this with your final patch.
> Or do you want me to fix the two bugs and resend ?
```

Sorry! I didn't have any time for the kernel hacking last days. There are some other minor (and not related) problems with the blocked signals which I'd like to check before doing the final patch.

Please feel free to do what you think right. I am going to KS this night, and I will be completely offline till september 10. Any chance we could delay this a bit? In any case, patches 2-3 should not depend on any further possible changes in this patch.

Oleg.

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