
Subject: Re: [PATCH 1/2] signals: kill(-1) should only signal processes in the same namespace

Posted by [Oleg Nesterov](#) on Thu, 17 Jul 2008 17:31:10 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 07/17, Pavel Emelyanov wrote:

>
> Daniel Hokka Zakrisson wrote:
> >
> > The way zap_pid_ns_processes does it is worse, since it signals every
> > thread in the namespace rather than every thread group. So either we walk
>
> It's questionable whether there are more "threads in a pid namespace" than
> "processes in a system".
>
> E.g. on my notebook there are ~110 processes and ~150 threads. So having
> this setup launched in 10 containers you'll have to walk 1100 tasks, while
> zap_pid_ns_processes only 150 ;)
>
> Some real-life example with containers: on one of our servers with 10
> containers serving as git repo, bulding system and some other stuff there
> are ~200 process totally and ~20 threads in each container. See?
>
> I tend to believe that walking threads in a container is cheaper then
> walking processes in a system...

kill_something_info() can't walk threads, think about the realtime signals.

Anyway, I think we should change kill_something_info(-1) to use rcu_read_lock()
instead of tasklist.

Oleg.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
