
Subject: Re: [PATCH 1/3] Signal semantics for /sbin/init
Posted by [Cedric Le Goater](#) on Thu, 13 Sep 2007 15:40:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

Oleg Nesterov wrote:

> On 09/10, sukadev@us.ibm.com wrote:

>> (This is Oleg's patch with my pid ns additions. Compiled and unit tested
>> on 2.6.23-rc4-mm1 with other patches in this set. Oleg pls update this
>> patch if necessary and sign-off)

>

> Sukadev, my apologies. This patch does need some changes,

>

>> Notes:

>>

>> - Blocked signals are never ignored, so init still can receive
>> a pending blocked signal after sigprocmask(SIG_UNBLOCK).
>> Easy to fix, but probably we can ignore this issue.

>

> I was wrong. This should be fixed right now. I think this is easy,
> and I was going to finish this patch yesterday, but - sorry! - I just
> can't switch to "kernel mode" these days, I am fighting with some urgent
> tasks on my paid job.

>

> Please feel free to solve this issue yourself if you wish. As for me, I'm
> forgetting about the kernel until at least the next weekend.

>

> Also, Pavel has some ideas how to do this all on receiver's path, perhaps
> this makes more sense (not that I completely agree, but I didn't see the
> code yet).

To respect the current init semantic, shouldn't we discard any unblockable
signal (STOP and KILL) sent by a process to its pid namespace init process ?
Then, all other signals should be handled appropriately by the pid namespace
init.

We are assuming that the pid namespace init is not doing anything silly and
I guess it's OK if the consequences are only on the its pid namespace and
not the whole system.

Cheers,

C.

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