
Subject: Re: [PATCH 3/15] kern_siginfo helper
Posted by [Pavel Emelianov](#) on Mon, 30 Jul 2007 06:07:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

Oleg Nesterov wrote:

> On 07/26, Pavel Emelianov wrote:
>> TODO: This is more an exploratory patch and modifies only interfaces
>> necessary to implement correct signal semantics in pid namespaces.
>>
>> If the approach is feasible, we could consistently use 'kern_siginfo'
>> in other signal interfaces and possibly in 'struct sigqueue'.
>>
>> We could modify dequeue_signal() to directly work with 'kern_siginfo'
>> and remove dequeue_signal_kern_info().
>
> Well... I know, it is very easy to blame somebody else's patch, and probably
> my taste is not good...
>
> But honestly, I personally think this approach is a horror, and any alternative
> is better :)
>
> I'd rather change dequeue_signal() so that it takes "struct sigqueue *"
> parameter instead of "siginfo_t *", or add a new "int *flags".
>
> OK, this doesn't work anyway, we should do something different. Perhaps
> just do all checks on sender's side.

Yes. Signal handling in namespaces turned out to be the most complicated part of the set. I start thinking to drop this part till we have the "core" in -mm tree. Suka, what do you think?

> It is a bit strange that this patch is 3/15, and the rest bits in 11/15,
> not very convenient for the review.

Well, I thought that a split like

1. patches for kernel to prepare it for the set
2. the set itself

is the best to review. Maybe I was wrong, but how to make this then?

E.g. I have a MS_KERNOUNT patch, but its changes are used *much* later.

> Oleg.
>
>

Thanks,
Pavel
