
Subject: Re: [PATCH] [RFC] network namespace locking rules
Posted by [Daniel Lezcano](#) on Fri, 28 Sep 2007 16:27:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Denis V. Lunev wrote:

>
> -----
>
> Subject:
> Re: [PATCH] [RFC] network namespace locking rules
> From:
> ebiederm@xmission.com (Eric W. Biederman)
> Date:
> Thu, 27 Sep 2007 02:45:34 -0600
> To:
> "Denis V. Lunev" <den@sw.ru>
>
> To:
> "Denis V. Lunev" <den@sw.ru>
> CC:
> Daniel Lezcano <dlezcano@fr.ibm.com>, Benjamin Thery
> <benjamin.thery@bull.net>, Kirill Korotaev <dev@sw.ru>
>
>
> "Denis V. Lunev" <den@sw.ru> writes:
>
>> Hello, Eric!
>>
>> Unfortunately, I was wrong that your patch is sane :(
>>
>> It breaks current RTNL socket semantic. Namely, current code relies that
>> - netlink from user-space is queued to RTNL socket if RTNL lock is held
>> - all pending messages in that queue will be processed in rtnl_unlock
>
> I know we come very close to this but I have a hard time seeing
> this being guaranteed. We don't hold a lock so I think it is
> possible for a new message to come in via another path on SMP,
> and we miss it in rtnl_unlock. Although missing that message
> from both paths that grabs rtnl_lock sounds unlikely.
>

Thanks, I missed this one.
it makes more sense now :)
