Subject: Re: [NETLINK] Don't attach callback to a going-away netlink socket
Posted by xemul on Wed, 18 Apr 2007 08:59:20 GMT
View Forum Message <> Reply to Message

Evgeniy Polyakov wrote:
> On Wed, Apr 18, 2007 at 12:32:40PM +0400, Pavel Emelianov (xemul@sw.ru) wrote:
>> Evgeniy Polyakov wrote:
>>> On Wed, Apr 18, 2007 at 12:16:18PM +0400, Pavel Emelianov (xemul@sw.ru) wrote:
>>>> Sorry, I forgot to put netdev and David in Cc when I first sent it.
>>>>
>>>> There is a race between netlink_dump_start() and netlink_release()
>>>> that can lead to the situation when a netlink socket with non-zero
>>>> callback is freed.
>>> Out of curiosity, why not to fix a netlink_dump_start() to remove
>>> callback in error path, since in 'no-error' path it removes it in
>> Error path is not relevant here. The problem is that we
>> keep a calback on a socket that is about to be freed.
>
> Yes, you are right, that it will not be freed in netlink_release(),
> but it will be freed in netlink_dump() after it is processed (in no-error
> path only though).
>

But error path will leak it. On success path we would have
a leaked packet in sk_write_queue, since we did't see it in
skb_queue_purge() while doing netlink_release().

Of course we can place the struts in code to handle the case
when we have a released socket with the attached callback, but
it is more correct (IMHO) not to allow to attach the callbacks
to dead sockets.