Subject: Re: [NETNS] Oops in register_pernet_operations() with CONFIG_NET_NS=n

Posted by ebiederm on Thu, 25 Oct 2007 16:39:41 GMT

View Forum Message <> Reply to Message

Benjamin Thery <benjamin.thery@bull.net> writes:

```
> Eric W. Biederman wrote:
>> Benjamin Thery <benjamin.thery@bull.net> writes:
>>
>>> Denis V. Lunev wrote:
>>>> The patch attached should help. The idea is simple. The "init" should be
>>>> called only once without NETNS. Period. No need for any lists.
>>> This is the kind of idea I had but I didn't think it could be
>>> that simple. :)
>>> Thanks Denis.
>>
>> It isn't.
>>
>>>> I'll resend it to Dave after the ACK.
>>> Tested on x86 64 with CONFIG NET NS=n and y.
>>> It fixes the issue we observed.
>>> Acked-by: Benjamin Thery <benjamin.thery@bull.net>
>> Try rmmod.
> rmmod was part of my tests and it does work.
> I did:
> $ iptables --list
>
   modules x_tables, ip_tables & iptable_filter are loaded
   each calling register_pernet_subsys.
>
> $ rmmod iptable filter ip tables x tables
   No problem here
> $ iptables --list
  To be sure I can load the modules again.
```

If that is the case the symptoms you are seeing make sense.

You haven't changed those modules to be mark struct

pernet_operations as __net_initdata have you?

Not doing the list walks helps when if it is only compiled in kernel data structures that are removed. However if it is potentially modular data structures that are removed the dereference of exit in unregister_pernet_subsys will also have problems.

_		
_	rı	\sim
ᆫ	11	U

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