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

Posted by [ebiederm](#) on Thu, 25 Oct 2007 17:21:55 GMT

[View Forum Message](#) <> [Reply to Message](#)

"Denis V. Lunev" <dlunev@gmail.com> writes:

> Eric W. Biederman wrote:

>> 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.

>

> this will work due to INIT_LIST_HEAD with circles list to itself and a

> del operation will work.

Suppose I have this fragment of code in a module:

```
> static int __net_init xt_net_init(struct net *net)
```

```
> {
```

```
>     ...
```

```
> }
```

```
>
```

```
> static void __net_exit xt_net_exit(struct net *net)
```

```
> {
```

```
>     ...
```

```
> }
```

```
>
```

```
> static struct pernet_operations __net_initdata xt_net_ops = {
```

```
>     .init = xt_net_init,
```

```
>     .exit = xt_net_exit,
```

```
> };
```

```
>
```

```
> static int __init xt_init(void)
```

```
> {
```

```
>     return register_pernet_subsys(&xt_net_ops);
```

```
> }
```

```
>
```

```
> static void __exit xt_fini(void)
```

```
> {
```

```
> unregister_pernet_subsys(&xt_net_ops);  
> }  
>  
> module_init(xt_init);  
> module_exit(xt_fini);
```

What happens during module removal when unregister_pernet_subsys calls xt_net_ops.exit after xt_net_ops has been removed from the kernel's memory?

```
> By the way, I think that we can in the case of undefined CONFIG_NET_NS  
> reduce register to calling ->init method and unregister to calling  
> ->exit method.  
>  
> This is a correct thing at least for now and will be welcomed by the all  
> embedded/etc people.
```

I'm not fundamentally opposed. Earlier versions of my patchset did that and more. However I think the pain is greater than the gain right now. Especially since this concept seems to require having quality inspected into it.

Eric

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