
Subject: Re: NET namespace locking seems broken to me
Posted by [ebiederm](#) on Fri, 21 Sep 2007 07:54:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

"Denis V. Lunev" <den@sw.ru> writes:

> Eric W. Biederman wrote:

>> Ok. I have found an important case. loopback.

> May be it will be better to move this in netdev_run_todo to cleanup
> locking. I am not sure right now.

? register_netdev when we create a new loopback device.

> Basically, there are 4 (four) locks after the patch:

> - dev_base_lock

> - rtnl

> - net_list_mutex

> - net_mutex

>

> Too many for me :)

I won't argue that. Especially as we don't have all the uses cases merged just yet. But after going through the reasoning for why I have them I don't see a good alternative. The only thing going for us is that it is very rare that something will want to traverse the list of network namespaces.

I'm guessing that something from this thread should become a big fat common on those locks. At the very least saying:

net_lock -> net_list_lock -> rtnl_lock.

> Why do we need them? The only case is that we want absence of some
> protocols/layers inside different namespaces. We have the only rtnl
> socket in OpenVZ

I may be wrong but I have a dim memory of your sockets not being per network namespace (or the equivalent) in OpenVZ. Which if my memory is correct probably accounts for most of the differences in implementation.

>> You want to cook up the patch to fix register_netdevice_notifier?

>

> I am trying this now.

Thanks.

