
Subject: NET namespace locking seems broken to me

Posted by [den](#) on Thu, 20 Sep 2007 11:29:14 GMT

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

Hello, Eric!

Current locking in mainstream seems broken to me.

1. struct net->list is manipulated under double net_mutex/net_list_mutex
2. net_list_mutex has been taken only in cleanup_net/net_ns_init inside net_mutes and seems pointless now
3. for_each_net (iterating against net_namespace_list) is called from
 - a) register_netdevice_notifier/__rtnl_link_unregister
 - b) register_pernet_operations/unregister_pernet_operationsIn the case b) the situation is sane, i.e. net_mutex is held while in the case b) we held rtnl_only

So, this does not look good to me for now.

How to cure this situation? I think that we can drop all locks for now and perform all operations under rtnl only. In the other case we must decide now should we make rtnl inner or outer for net_mutex.

>From my point of view net_mutex should be taken inside rtnl lock and we must add it now into list manipulation routines.

Plz point me to my mistake in logic :)

Regards,
Den

P.S. The GIT with your latest tree is also affected
