
Subject: Re: [PATCH] [NETNS45] network namespace locking rules

Posted by [ebiederm](#) on Fri, 28 Sep 2007 16:54:41 GMT

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

"Denis V. Lunev" <den@openvz.org> writes:

> Current locking for network namespace list/initialization is broken.
> for_each_net is called under single rtnl_lock in
> register_netdevice_notifier.

As of 984e617a3e1974022b8f671427a76ffbe886f75b this issue has been addressed in net-2.6.24

The only remaining part to address is rtnl_unlock().

My current hypothesis is that rtnl_unlock() only needs to process the packets that are queued while we had the rtnl_lock held, to maintain the current semantics.

So the retry may not be necessary as it may be possible to prove that the only extra packets that could come in come from another thread taking the rtnl_lock and they will those packets in their rtnl_unlock() if we don't.

I will review this later today, and add the retry if necessary.

One way or another I think we agree on how to get the locking correct. The other details are a bit tricky but look usable.

Eric
