
Subject: Re: [PATCH 2/5] net: Make rtnetlink infrastructure network namespace aware

Posted by [Patrick McHardy](#) on Sat, 29 Sep 2007 17:48:03 GMT

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

Eric W. Biederman wrote:

> Patrick McHardy <kaber@trash.net> writes:

>

>

>>I'm wondering why this receive queue processing on unlock is still

>>necessary today, we don't do trylock in rtnetlink_rcv anymore, so

>>all senders will simply wait until the lock is released and then

>>process the queue.

>

>

> Good question, I should probably look. I was lazy and didn't go back

> and audit why we were doing this. I just coded a routine that I was

> certain would work. It does appear that we are processing the queue

> with sk_data_read when we add a message, so this may be completely

> unnecessary. I will go back and look. If we can remove this bit

> things should be simpler.

Maybe I can save you some time: we used to do down_trylock() for the rtnl mutex, so senders would simply return if someone else was already processing the queue *or* the rtnl was locked for some other reason. In the first case the process already processing the queue would also process the new messages, but if it the rtnl was locked for some other reason (for example during module registration) the message would sit in the queue until the next rtnetlink sendmsg call, which is why rtnl_unlock does queue processing. Commit 6756ae4b changed the down_trylock to mutex_lock, so senders will now simply wait until the mutex is released and then call netlink_run_queue themselves. This means its not needed anymore.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
