
Subject: Re: [PATCH v3] tun: don't hold network namespace by tun sockets

Posted by [davem](#) on Tue, 13 Mar 2012 00:14:11 GMT

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

aFrom: Stanislav Kinsbursky <skinsbursky@parallels.com>

Date: Mon, 12 Mar 2012 16:59:41 +0400

> v3: added previously removed sock_put() to the tun_release() callback, because
> sk_release_kernel() doesn't drop the socket reference.
>
> v2: sk_release_kernel() used for socket release. Dummy tun_release() is
> required for sk_release_kernel() ---> sock_release() ---> sock->ops->release()
> call.
>
> TUN was designed to destroy it's socket on network namespace shutdown. But this
> will never happen for persistent device, because it's socket holds network
> namespace.
> This patch removes of holding network namespace by TUN socket and replaces it
> by creating socket in init_net and then changing it's net to desired one. On
> shutdown socket is moved back to init_net prior to final put.
>
> Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

Looks good, applied, thanks.
