

---

Subject: [PATCH v2] tun: don't hold network namespace by tun sockets

Posted by [Stanislav Kinsbursky](#) on Mon, 12 Mar 2012 11:20:19 GMT

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

---

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 it to the desired one. On shutdown socket is moved back to init\_net prior to final put.

Signed-off-by: Stanislav Kinsbursky <[skinsbursky@parallels.com](mailto:skinsbursky@parallels.com)>

---

drivers/net/tun.c | 13 ++++++++  
1 files changed, 10 insertions(+), 3 deletions(-)

diff --git a/drivers/net/tun.c b/drivers/net/tun.c

index 2c5d349..b13ee48 100644

--- a/drivers/net/tun.c

+++ b/drivers/net/tun.c

@@ -359,7 +359,7 @@ static void tun\_free\_netdev(struct net\_device \*dev)

{  
 struct tun\_struct \*tun = netdev\_priv(dev);

- sock\_put(tun->socket.sk);  
+ sk\_release\_kernel(tun->socket.sk);  
}

/\* Net device open. \*/

@@ -980,10 +980,16 @@ static int tun\_rcvmsg(struct kiocb \*iocb, struct socket \*sock,  
 return ret;  
}

+static int tun\_release(struct socket \*sock)

+{  
+ return 0;  
+}

+

/\* Ops structure to mimic raw sockets with tun \*/

static const struct proto\_ops tun\_socket\_ops = {  
 .sendmsg = tun\_sendmsg,  
 .rcvmsg = tun\_rcvmsg,  
+ .release = tun\_release,

```

};

static struct proto tun_proto = {
@@ -1110,10 +1116,11 @@ static int tun_set_iff(struct net *net, struct file *file, struct ifreq *ifr)
    tun->vnet_hdr_sz = sizeof(struct virtio_net_hdr);

    err = -ENOMEM;
-   sk = sk_alloc(net, AF_UNSPEC, GFP_KERNEL, &tun_proto);
+   sk = sk_alloc(&init_net, AF_UNSPEC, GFP_KERNEL, &tun_proto);
    if (!sk)
        goto err_free_dev;

+   sk_change_net(sk, net);
    tun->socket.wq = &tun->wq;
    init_waitqueue_head(&tun->wq.wait);
    tun->socket.ops = &tun_socket_ops;
@@ -1174,7 +1181,7 @@ static int tun_set_iff(struct net *net, struct file *file, struct ifreq *ifr)
    return 0;

err_free_sk:
-   sock_put(sk);
+   tun_free_netdev(dev);
err_free_dev:
    free_netdev(dev);
failed:

```

---

Subject: Re: [PATCH v2] tun: don't hold network namespace by tun sockets  
 Posted by [Stanislav Kinsbursky](#) on Mon, 12 Mar 2012 13:02:54 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Drop this patch. There will be third version.  
 Sorry for inconvenience.

--  
 Best regards,  
 Stanislav Kinsbursky

---

Subject: Re: [PATCH v2] tun: don't hold network namespace by tun sockets  
 Posted by [Stanislav Kinsbursky](#) on Mon, 12 Mar 2012 13:05:31 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Drop this patch. There will be third version.  
 Sorry for inconvenience.

--

Best regards,  
Stanislav Kinsbursky

---

---

Subject: Re: [PATCH v2] tun: don't hold network namespace by tun sockets  
Posted by [Stanislav Kinsbursky](#) on Mon, 12 Mar 2012 13:07:35 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Drop this patch. There will be third version.  
Sorry for inconvenience.

--

Best regards,  
Stanislav Kinsbursky

---