
Subject: Re: [PATCH net-2.6.26 3/3][TUN][NETNS]: Allow to register tun devices in namespace.

Posted by [Pavel Emelianov](#) on Thu, 03 Apr 2008 12:11:23 GMT

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

Max Krasnyanskiy wrote:

> Hi Pavel,

>

> All three patches look perfectly fine to me.

> Feel free to add Acked-by: Max Krasnyansky <maxk@qualcomm.com>

:)

> Pavel Emelianov wrote:

>> This is basically means that a net is set for a new device, but

>> actually this involves two more steps:

>>

>> 1. mark the tun device as "local", i.e. do not allow for it to

>> move across namespaces.

>>

>> This is done so, since tun device is most often associated to some

>> file (and thus to some process) and moving the device alone is not

>> valid while keeping the file and the process outside.

>>

>> 2. get the tun device's net when tun becomes attached and put one

>> when it becomes detached.

>>

>> This is needed to handle the case when a task owning the tun dies,

>> but a files lives for some more time - in this case we must not

>> allow for net to be freed, since its exit hook will spoil that file's

>> private data by unregistering the tun from under tun_chr_close.

> I'm not sure what you mean "by file lives on" here. I believe you're talking

I meant the following. When a task dies it releases it's net namespace and thus this net cleanup job may start. But the file, that holds the (not yet detached) tun device is not released immediately - so before this file is being released the tun_exit_net callback may be called and the file's tun device will be unregistered in it. I catch this race by getting the net by attached tun device, so that net is cleaned up only when there are only detached tuns in it. In this case no tasks will pick up these tuns and everything will be OK.

> about persistent tun devices. ie Those that exist in detached state and are
> not attached to any file descriptors.

> "net" refcounting logic there looks fine to me.

Great!

> Max

Thanks,
Pavel
