
Subject: Re: [PATCH 10/12] L2 network namespace: playing with pass-through device

Posted by [Daniel Lezcano](#) on Tue, 12 Dec 2006 14:10:21 GMT

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

Dmitry Mishin wrote:

> On Tuesday 12 December 2006 16:59, Daniel Lezcano wrote:

>> Dmitry Mishin wrote:

>>> Temporary code to debug and play with pass-through device.

>>> Create device pair by

>>> modprobe veth

>>> echo 'add veth1 0:1:2:3:4:1 eth0 0:1:2:3:4:2' >/proc/net/veth_ctl

>>> and your shell will appear into a new namespace with `eth0' device.

>>> Configure device in this namespace

>>> ip l s eth0 up

>>> ip a a 1.2.3.4/24 dev eth0

>>> and in the root namespace

>>> ip l s veth1 up

>>> ip a a 1.2.3.1/24 dev veth1

>>> to establish a communication channel between root namespace and the newly

>>> created one.

>>>

>>> Code is done by Andrey Savochkin and ported by me over Cedric's patchset

>>>

>>> Signed-off-by: Dmitry Mishin <dim@openvz.org>

>>>

>> [...]

>>

>>> --- linux-2.6.19-rc6-mm2.orig/include/linux/net_namespace.h

>>> +++ linux-2.6.19-rc6-mm2/include/linux/net_namespace.h

>>> @@ -24,6 +24,9 @@ struct net_namespace {

>>> int fib4_tribe_last_dflt;

>>> #endif

>>> unsigned int hash;

>>> + struct net_namespace *parent;

>>> + struct list_head child_list, sibling_list;

>>> + unsigned int id;

>>> };

>> Why do you need to have a child list and sibling list ?

> Because of the level2<->level3 hierarchy, for example.

This hierarchy doesn't exist with ns->parent ? Do you have an example when the hierarchy should be used ? I mean when we need to browse from I2 -> I3 ?
