

Hi, David.

What are your plans about this driver?

Thanks,
Pavel

- > Veth stands for Virtual ETHernet. It is a simple tunnel driver
- > that works at the link layer and looks like a pair of ethernet
- > devices interconnected with each other.
- >
- > Mainly it allows to communicate between network namespaces but
- > it can be used as is as well. E.g. one may join to bridged
- > networkings segments together.
- >
- > Eric recently sent a similar driver called etun with the sysfs
- > interface. This implementation uses another interface - the RTM_NRELINK
- > message introduced by Patric.
- >
- > The newlink callback is organized that way to make it easy to create the
- > peer device in the separate namespace when we have them in kernel.
- >
- > Many thanks to Patrick for reviewing the patches and his advises
- > on how to make driver cleaner.
- >
- > Changes from v.3:
- > * Reserved place for struct ifinfomsg in IFLA_INFO_DATA part
- > of the packet. This is not used yet, but may be in the
- > future.
- >
- > Changes from v.2.1:
- > * Made the generic routine for link creation to be used
- > by veth driver, any other tunnel driver that needs to
- > create several devices at once and rtnl_newlink() code.
- >
- > Changes from v.2:
- > * Rebase over latest netdev tree. No actual changes;
- > * Small code rework.
- >
- > Changes from v.1:
- > * Per-cpu statistics;
- > * Standard convention for nla policy names;
- > * Module alias added;
- > * Xmit function fixes noticed by Patric;

> * Code cleanup.
>
> The patch for an ip utility is also provided.
>
> Signed-off-by: Pavel Emelianov <xemul@openvz.org>
> Acked-by: Patrick McHardy <kaber@trash.net>
>
