
Subject: Re: [patch 0/1] [RFC][net namespace] veth ioctl management
Posted by [ebiederm](#) on Mon, 19 Feb 2007 14:29:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

Daniel Lezcano <dlezcano@fr.ibm.com> writes:

- > The following patch is an upgrade proposition for the veth pass-through driver.
- > The temporary proc interface has been replaced by an ioctl. The device is
- > a misc device. The major number is 10 and the minor number is dynamically
- > allocated.
- > The minor number should to be retrieved from the /dev/misc

Ugh.

I don't think it's perfect but please take a look at the sysfs interface to my etun device. I think it is a little easier to work with.

- > The veth configuration is no more done in one shot and does not unshare anymore.
- > The veth creation is done after the network namespace has been unshared and
- > binded with the bind ns.
- > The ioctl takes the parent dev name, the child dev name and the child namespace
- > identifier.
- > The veth creation will fails if the namespaces relationship is not
- > parent->child.

Although with a lot of what you describe veth sounds decidedly less flexible than etun.

- > The set_mac_address callback is set, so the usual commands ip/ifconfig can be
- > used to
- > assign a mac address to the pair device.
- >
- > The patch header contains the source code for the userspace tool, vethctl.
- > The patch needs the lxc patchset for the bind namespace syscall.

Eric

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>

Subject: Re: [patch 0/1] [RFC][net namespace] veth ioctl management
Posted by [Daniel Lezcano](#) on Mon, 19 Feb 2007 15:32:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> Daniel Lezcano <dlezcano@fr.ibm.com> writes:
>
>> The following patch is an upgrade proposition for the veth pass-through driver.
>> The temporary proc interface has been replaced by an ioctl. The device is
>> a misc device. The major number is 10 and the minor number is dynamically
>> allocated.
>> The minor number should to be retrieved from the /dev/misc
>
> Ugh.
>
> I don't think it's perfect but please take a look at the sysfs
> interface to my etun device. I think it is a little easier to
> work with.

I will have look to *your* etun interface. But, PLEASE, look at the net namespace network patches (L2/L3) which were rfc'ed, posted several months ago.

>
>> The veth configuration is no more done in one shot and does not unshare anymore.
>> The veth creation is done after the network namespace has been unshared and
>> binded with the bind ns.
>> The ioctl takes the parent dev name, the child dev name and the child namespace
>> identifier.
>> The veth creation will fails if the namespaces relationship is not
>> parent->child.
>
> Although with a lot of what you describe veth sounds decidedly less flexible
> than etun.

Unfortunately, you didn't say that when Andrey sent it this summer and when Dmitry ported it to the namespaces and when it was integrated to lxc patchset by Cedric and when I put the L3 namespace on top of it.

Eric, opensource is about collaboration.

-- Daniel

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
