
Subject: Re: [PATCH] net: Add etun driver
Posted by [Ben Greear](#) on Fri, 06 Apr 2007 21:20:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> etun is a simple two headed tunnel driver that at the link layer looks
> like ethernet. It's target audience is communicating between network
> namespaces but it is general enough it has other valid uses as well.
>
> Ben Greear implemented a similar device called redir-dev, for network
> emulation.
>
> OpenVZ has a similar device that goes by the name veth.
>
> I didn't want to mess with ioctls or weird non-general network
> interfaces for creating devices, so I used sysfs as my control
> mechanism.
>
> To create a pair of devices called veth0 and veth1:
> echo -n 'veth0,veth1' > /sys/module/etun/parameters/newif
>
> To destroy a pair of devices:
> echo -n 'veth0' > /sys/module/etun/parameters/delif
>

Is there any way to tell for certain if an interface is a etun or not?
Maybe
a file could be found (or not) in sysfs somewhere?

Also, how do you find the peer device from user-space? This would be
very useful
for anyone trying to manage these devices with a user-space program.

When you are creating new devices, I think you should check to make
sure there isn't already a device with that name.

In general though, I look forward to this being in the kernel so I can drop
my redirect device code from my out-of-tree patch.

Thanks,
Ben

--
Ben Greear <greearb@candelatech.com>
Candela Technologies Inc <http://www.candelatech.com>

Subject: Re: [PATCH] net: Add etun driver
Posted by [ebiederm](#) on Sat, 07 Apr 2007 02:06:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ben Greear <greearb@candelatech.com> writes:

> Is there any way to tell for certain if an interface is a etun or not? Maybe
> a file could be found (or not) in sysfs somewhere?

Link for any decent network driver `ethtool -i <interface>`

> Also, how do you find the peer device from user-space? This would be very
> useful
> for anyone trying to manage these devices with a user-space program.

Currently "`ethtool -S <interface>`"
And read the `partner_ifindex`.

Further whoever generates the pair specifies the initial set of names.

> When you are creating new devices, I think you should check to make
> sure there isn't already a device with that name.

`alloc_netdev` doesn't let you register two devices with the same name,
and I check the return status of `alloc_netdev`.

> In general though, I look forward to this being in the kernel so I can drop
> my redirect device code from my out-of-tree patch.

Thanks.

Eric

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [PATCH] net: Add etun driver
Posted by [Ben Greear](#) on Sat, 07 Apr 2007 03:31:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> Ben Greear <greearb@candelatech.com> writes:

>

>

>> Is there any way to tell for certain if an interface is a etun or not? Maybe

>> a file could be found (or not) in sysfs somewhere?

>>

>

> Link for any decent network driver ethtool -i <interface>

>

I guess that will do, but then if you ever change the strings, any

user-space that is

depending on this will break or have to be modified with additional

craft. It seems

cleaner to me to have an ioctl or a specific place in /proc or some

other virtual

fs, but I can deal with it either way...

>> Also, how do you find the peer device from user-space? This would be very

>> useful

>> for anyone trying to manage these devices with a user-space program.

>>

>

> Currently "ethtool -S <interface>"

> And read the partner_ifindex.

>

Ok, that will work. Again, my personal preference is for a single

specific ioctl or proc'ish file

to read the specific value instead of having to parse strings, but this

will do.

> Further whoever generates the pair specifies the initial set of names.

>

Yeah, but you can't depend on knowing that in an interesting environment.

Thanks,

Ben

--

Ben Greear <greearb@candelatech.com>

Candela Technologies Inc <http://www.candelatech.com>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [PATCH] net: Add etun driver
Posted by [ebiederm](#) on Sat, 07 Apr 2007 05:37:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ben Greear <greearb@candelatech.com> writes:

> I guess that will do, but then if you ever change the strings, any user-space
> that is
> depending on this will break or have to be modified with additional cruft. It
> seems
> cleaner to me to have an ioctl or a specific place in /proc or some other
> virtual
> fs, but I can deal with it either way...

True if the name of the driver changes from etun there is an issue.

>>> Also, how do you find the peer device from user-space? This would be very
>>> useful
>>> for anyone trying to manage these devices with a user-space program.
>>>
>>
>> Currently "ethtool -S <interface>"
>> And read the partner_ifindex.
>>
> Ok, that will work. Again, my personal preference is for a single specific
> ioctl or proc'ish file
> to read the specific value instead of having to parse strings, but this will do.

Hmm. I guess there is string parsing to identify the index.

I guess a sysfs device attribute would work as well.

>> Further whoever generates the pair specifies the initial set of names.
>>
> Yeah, but you can't depend on knowing that in an interesting environment.

Frankly. In an interesting environment I haven't been able to think of a way to successfully say anything about the partner device.

The problem is that all identifiers are namespace local so the remote side is not in the current namespace the ifindex or the device name mean nothing.

In that case the only remotely usable value I can return is the mac address of the other side.

Eric

Containers mailing list
Containers@lists.linux-foundation.org

Subject: Re: [PATCH] net: Add etun driver
Posted by [Ben Greear](#) on Sat, 07 Apr 2007 07:31:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> Ben Greear <greearb@candelatech.com> writes:

>

>

>> I guess that will do, but then if you ever change the strings, any user-space

>> that is

>> depending on this will break or have to be modified with additional cruft. It

>> seems

>> cleaner to me to have an ioctl or a specific place in /proc or some other

>> virtual

>> fs, but I can deal with it either way...

>>

>

> True if the name of the driver changes from etun there is an issue.

>

So, how about a sysfs field for this too, something like 'is-etun' or
whatever...

Or, IOCTL works fine too, of course.

>

>>>> Also, how do you find the peer device from user-space? This would be very

>>>> useful

>>>> for anyone trying to manage these devices with a user-space program.

>>>>

>>>>

>>> Currently "ethtool -S <interface>"

>>> And read the partner_ifindex.

>>>

>>>

>> Ok, that will work. Again, my personal preference is for a single specific

>> ioctl or proc'ish file

>> to read the specific value instead of having to parse strings, but this will do.

>>

>

> Hmm. I guess there is string parsing to identify the index.

>

> I guess a sysfs device attribute would work as well.

>

Yes, that would be much appreciated.

>

>>> Further whoever generates the pair specifies the initial set of names.

>>>

>>>

>> Yeah, but you can't depend on knowing that in an interesting environment.

>>

>

> Frankly. In an interesting environment I haven't been able to think of a

> way to successfully say anything about the partner device.

>

> The problem is that all identifiers are namespace local so the remote side

> is not in the current namespace the ifindex or the device name mean nothing.

>

> In that case the only remotely usable value I can return is the mac address

> of the other side.

>

Couldn't you also show the peer's name-space id in that case?

MAC could be duplicated, so that's not a very good key.

Thanks,

Ben

--

Ben Greear <greearb@candelatech.com>

Candela Technologies Inc <http://www.candelatech.com>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
