

---

Subject: Re: [patch 0/1] [RFC][net namespace] veth ioctl management  
Posted by [ebiederm](#) on Mon, 19 Feb 2007 17:30:41 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

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

> 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.

I have, and I have commented.

> Unfortunatly, 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.

Mostly it is six of one half a dozen of the other as far as the actual implementation is concerned. The practical difference is etun is not tied in any way shape or form to namespaces, whereas veth appears to be.

> Eric, opensource is about collaboration.

Yes. Exactly. That is why I am suggesting that we try to use something simpler than ioctl.

The process is about collaborating to find the best technical solution we can find in a timely manner.

Currently the L2 stuff derived from OpenVZ appears to be completely unmergable because of how the patchset is constructed.

I am very frustrated that you are attacking me about not cooperating. I have made comments (that were hopefully constructive) all along the way. At some points in time I was busy so I could not look at things in detail, but I have tried across the whole spectrum of the namespaces effort. I finally sat down and wrote my own network namespace implementation because some very important points were not getting addressed and I had some novel ideas on how to solve some of the problems. So I figured the contrast would be useful. The containers list wasn't my primary audience of my RFC but I was surprised that I received no comments there.

Plus this is the biggest problem I see with the OpenVZ derived network namespace effort is lack of cooperation with the people who matter. The maintainers of the network stack. Sure people who are actually going to use a network namespace have to agree that the code will solve the problem in a usable way but if the maintainers of the network stack don't like or can't be interested the code is never going anywhere.

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 [Mishin Dmitry](#) on Mon, 19 Feb 2007 17:36:59 GMT

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

---

On Monday 19 February 2007 20:30, Eric W. Biederman wrote:

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

> > 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.

>

> I have, and I have commented.

>

> > Unfortunatly, 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.

>

> Mostly it is six of one half a dozen of the other as far as the actual  
> implementation is concerned. The practical difference is etun is not tied  
> in any way shape or form to namespaces, whereas veth appears to be.

>

> > Eric, opensource is about collaboration.

>

> Yes. Exactly. That is why I am suggesting that we try to use something  
> simpler than ioctl.

>

> The process is about collaborating to find the best technical solution  
> we can find in a timely manner.

>

> Currently the L2 stuff derived from OpenVZ appears to be completely  
> unmergeable because of how the patchset is constructed.

Please, explain. For me, your code is not mergeable as well.

>

> I am very frustrated that you are attacking me about not cooperating.  
> I have made comments (that were hopefully constructive) all along the  
> way. At some points in time I was busy so I could not look at things  
> in detail, but I have tried across the whole spectrum of the  
> namespaces effort. I finally sat down and wrote my own network  
> namespace implementation because some very important points were not  
> getting addressed and I had some novel ideas on how to solve some of

> the problems. So I figured the contrast would be useful. The  
> containers list wasn't my primary audience of my RFC but I was  
> surprised that I received no comments there.  
>  
> Plus this is the biggest problem I see with the OpenVZ derived network  
> namespace effort is lack of cooperation with the people who matter.  
> The maintainers of the network stack. Sure people who are actually  
> going to use a network namespace have to agree that the code will  
> solve the problem in a usable way but if the maintainers of the  
> network stack don't like or can't be interested the code is never  
> going anywhere.  
Fully agree. But as I can see, your code arises no more comments, than ours.  
So, we need to find other ways. Do you have more ideas?

--

Thanks,  
Dmitry.

---

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 [dev](#) on Mon, 19 Feb 2007 18:24:31 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Eric,

> Mostly it is six of one half a dozen of the other as far as the actual implementation  
> is concerned. The practical difference is etun is not tied in any way shape or  
> form to namespaces, whereas veth appears to be.  
veth is not tied to namespaces as well.  
it is possible to create a pair of veth devices in host system for example.

> Currently the L2 stuff derived from OpenVZ appears to be completely  
> unmergeable because of how the patchset is constructed.  
you didn't say it before... :)  
can you point to \*exact\* issues we can fix?  
I really think that commenting each other approaches  
and fixing this we can go quicker.

> I am very frustrated that you are attacking me about not cooperating.  
> I have made comments (that were hopefully constructive) all along the  
> way. At some points in time I was busy so I could not look at things  
> in detail, but I have tried across the whole spectrum of the  
> namespaces effort. I finally sat down and wrote my own network

> namespace implementation because some very important points were not  
> getting addressed and I had some novel ideas on how to solve some of  
> the problems. So I figured the contrast would be useful. The  
> containers list wasn't my primary audience of my RFC but I was  
> surprised that I received no comments there.

> Plus this is the biggest problem I see with the OpenVZ derived network  
> namespace effort is lack of cooperation with the people who matter.  
> The maintainers of the network stack. Sure people who are actually  
> going to use a network namespace have to agree that the code will  
> solve the problem in a usable way but if the maintainers of the  
> network stack don't like or can't be interested the code is never  
> going anywhere.

I hope David is going to visit OLS so we could discuss it orally.

Until that I will try to ask about his position in the email.

Thanks,

Kirill

P.S. as for cooperating - we probably all do it in a poor manner yet,  
at least this is my personal feeling.

---

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 Tue, 20 Feb 2007 09:24:39 GMT

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

---

Eric W. Biederman wrote:

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

>

>> 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.

>

> I have, and I have commented.

>

>> Unfortunatly, 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.

>

> Mostly it is six of one half a dozen of the other as far as the actual implementation  
> is concerned. The practical difference is etun is not tied in any way shape or  
> form to namespaces, whereas veth appears to be.

>

>> Eric, opensource is about collaboration.

>  
> Yes. Exactly. That is why I am suggesting that we try to use something  
> simpler than ioctl.

Why, don't you send a patch with your etun driver on top of the network namespace patches ? So we will be able to try it ...

> The process is about collaborating to find the best technical solution  
> we can find in a timely manner.  
>  
> Currently the L2 stuff derived from OpenVZ appears to be completely  
> unmergeable because of how the patchset is constructed.  
>  
> I am very frustrated that you are attacking me about not cooperating.  
> I have made comments (that were hopefully constructive) all along the  
> way. At some points in time I was busy so I could not look at things  
> in detail, but I have tried across the whole spectrum of the  
> namespaces effort. I finally sat down and wrote my own network  
> namespace implementation because some very important points were not  
> getting addressed

Which ones ?

Why did you directly code before discussing with us about these points ?

>and I had some novel ideas on how to solve some of  
> the problems. So I figured the contrast would be useful. The  
> containers list wasn't my primary audience of my RFC but I was  
> surprised that I received no comments there.

Probably because you send directly a \*big\* patch touching a lot of file.  
Personally, I was quite surprise to discover you were working on your  
own network namespace while we were all working together to find a  
common solution L2/L3. This is an individualist behavior.

> Plus this is the biggest problem I see with the OpenVZ derived network  
> namespace effort is lack of cooperation with the people who matter.  
> The maintainers of the network stack. Sure people who are actually  
> going to use a network namespace have to agree that the code will  
> solve the problem in a usable way but if the maintainers of the  
> network stack don't like or can't be interested the code is never  
> going anywhere.

Sure.

---

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

---