
Subject: 2nd interface loses packets on using veth
Posted by Julian Yap **on Thu, 15 Mar 2007 03:16:43 GMT**
[View Forum Message](#) <> [Reply to Message](#)

Kind of related to my previous posting.

I've now setup my 2 VE interfaces using veth method.

No longer using Source based routing, I've set up my default route to be 192.168.0.1.

Pinging on the 192.168.0.x subnet is perfect from the VE.

Pinging on my VE's 2nd interface (eth1) works OK in burst but at times just drops packets.

Running a tcpdump, I see a log of these messages:

17:10:29.813495 arp who-has machine15 tell 192.168.1.247

When it does work OK, it's because it's returned this:

17:14:35.829082 arp who-has machine15 tell 192.168.100.247

17:14:36.828053 arp who-has machine13 tell machine15

...

17:14:36.828082 arp reply machine13 is-at 00:16:3e:7e:a5:2e

Any clues why it arp's so much?

- Julian

Subject: Re: 2nd interface loses packets on using veth
Posted by Andrey Mirkin **on Thu, 15 Mar 2007 07:39:05 GMT**
[View Forum Message](#) <> [Reply to Message](#)

Hello,

Please provide ip configuration (ip a ls) and route tables (ip r ls) for host system and VE.

Regards,
Andrey

On Thursday 15 March 2007 06:16 Julian Yap wrote:

> Kind of related to my previous posting.
>
> I've now setup my 2 VE interfaces using veth method.
>
> No longer using Source based routing, I've set up my default route to

> be 192.168.0.1.
>
> Pinging on the 192.168.0.x subnet is perfect from the VE.
>
> Pinging on my VE's 2nd interface (eth1) works OK in burst but at times
> just drops packets.
>
> Running a tcpdump, I see a log of these messages:
> 17:10:29.813495 arp who-has machine15 tell 192.168.1.247
>
> When it does work OK, it's because it's returned this:
> 17:14:35.829082 arp who-has machine15 tell 192.168.100.247
> 17:14:36.828053 arp who-has machine13 tell machine15
> ...
> 17:14:36.828082 arp reply machine13 is-at 00:16:3e:7e:a5:2e
>
> Any clues why it arp's so much?
>
> - Julian

Subject: Re: 2nd interface loses packets on using veth
Posted by [Julian Yap](#) on Thu, 15 Mar 2007 18:16:59 GMT

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

Thanks.

HOST: ip a ls

```
2: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        inet6 ::1/128 scope host
            valid_lft forever preferred_lft forever
4: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
6: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:18:8b:77:f2:01 brd ff:ff:ff:ff:ff:ff
    inet 192.168.100.247/24 brd 192.168.100.255 scope global eth0
        inet6 fe80::218:8bff:fe77:f201/64 scope link
            valid_lft forever preferred_lft forever
8: eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:18:8b:77:f2:02 brd ff:ff:ff:ff:ff:ff
    inet 10.100.225.247/24 brd 10.100.225.255 scope global eth1
        inet6 fe80::218:8bff:fe77:f202/64 scope link
            valid_lft forever preferred_lft forever
1: venet0: <BROADCAST,POINTOPOINT,NOARP> mtu 1500 qdisc noqueue
    link/void
```

```
3: veth101.0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:16:3e:23:d5:dd brd ff:ff:ff:ff:ff:ff
    inet6 fe80::216:3eff:fe23:d5dd/64 scope link
        valid_lft forever preferred_lft forever
5: veth101.1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:16:3e:7e:a5:2e brd ff:ff:ff:ff:ff:ff
    inet6 fe80::216:3eff:fe7e:a52e/64 scope link
        valid_lft forever preferred_lft forever
```

HOST: ip r ls

```
10.100.225.106 dev veth101.1 scope link
192.168.100.106 dev veth101.0 scope link
192.168.100.0/24 dev eth0 proto kernel scope link src 192.168.100.247
10.100.225.0/24 dev eth1 proto kernel scope link src 10.100.225.247
169.254.0.0/16 dev eth1 scope link
default via 192.168.100.1 dev eth0
```

VE: ip a ls

```
1: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
3: venet0: <BROADCAST,POINTOPOINT,NOARP> mtu 1500 qdisc noqueue
    link/void
5: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:16:3e:69:ce:df brd ff:ff:ff:ff:ff:ff
    inet 192.168.100.106/32 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::216:3eff:fe69:cедf/64 scope link
        valid_lft forever preferred_lft forever
7: eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:16:3e:1d:0d:54 brd ff:ff:ff:ff:ff:ff
    inet 10.100.225.106/32 scope global eth1
        valid_lft forever preferred_lft forever
    inet6 fe80::216:3eff:fe1d:d54/64 scope link
        valid_lft forever preferred_lft forever
```

VE: ip r ls

```
192.168.100.0/24 dev eth0 scope link
10.100.225.0/24 dev eth1 scope link
default dev eth0 scope link
```

On 3/14/07, Andrey Mirkin <major@openvz.org> wrote:

> Hello,
>
> Please provide ip configuration (ip a ls) and route tables (ip r ls) for host

> system and VE.
>
> Regards,
> Andrey
>
> On Thursday 15 March 2007 06:16 Julian Yap wrote:
> > Kind of related to my previous posting.
> >
> > I've now setup my 2 VE interfaces using veth method.
> >
> > No longer using Source based routing, I've set up my default route to
> > be 192.168.0.1.
> >
> > Pinging on the 192.168.0.x subnet is perfect from the VE.
> >
> > Pinging on my VE's 2nd interface (eth1) works OK in burst but at times
> > just drops packets.
> >
> > Running a tcpdump, I see a log of these messages:
> > 17:10:29.813495 arp who-has machine15 tell 192.168.1.247
> >
> > When it does work OK, it's because it's returned this:
> > 17:14:35.829082 arp who-has machine15 tell 192.168.100.247
> > 17:14:36.828053 arp who-has machine13 tell machine15
> > ...
> > 17:14:36.828082 arp reply machine13 is-at 00:16:3e:7e:a5:2e
> >
> > Any clues why it arp's so much?
> >
> > - Julian

Subject: Re: [SOLVED] 2nd interface loses packets on using veth
Posted by [Julian Yap](#) on Thu, 15 Mar 2007 23:37:23 GMT

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

Hi guys.

I moved back to using Venet since from the VE 'ping -I 10.100.225.106 10.100.225.1' worked for me. I figured maybe it was an ARP issue and could be sorted out with the "Multiple Network Interfaces And ARP Flux" wiki page instructions. No luck when I tried that.

I took another look in the forums and found my solution.

Running something like this from my VE works beautifully:
ip route add 10.100.225.0/24 dev venet0 src 10.100.225.106

Everything I've read so far suggests using Veth or Source based routing. Both of which didn't work for me.

- Julian

On 3/15/07, Julian Yap <julianokyap@gmail.com> wrote:

```
> Thanks.  
>  
> HOST: ip a ls  
>  
> 2: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue  
>     link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00  
>     inet 127.0.0.1/8 scope host lo  
>     inet6 ::1/128 scope host  
        valid_lft forever preferred_lft forever  
> 4: sit0: <NOARP> mtu 1480 qdisc noop  
>     link/sit 0.0.0.0 brd 0.0.0.0  
> 6: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000  
>     link/ether 00:18:8b:77:f2:01 brd ff:ff:ff:ff:ff:ff  
>     inet 192.168.100.247/24 brd 192.168.100.255 scope global eth0  
>     inet6 fe80::218:8bff:fe77:f201/64 scope link  
        valid_lft forever preferred_lft forever  
> 8: eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000  
>     link/ether 00:18:8b:77:f2:02 brd ff:ff:ff:ff:ff:ff  
>     inet 10.100.225.247/24 brd 10.100.225.255 scope global eth1  
>     inet6 fe80::218:8bff:fe77:f202/64 scope link  
        valid_lft forever preferred_lft forever  
> 1: venet0: <BROADCAST,POINTOPOINT,NOARP> mtu 1500 qdisc noqueue  
>     link/void  
> 3: veth101.0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue  
>     link/ether 00:16:3e:23:d5:dd brd ff:ff:ff:ff:ff:ff  
>     inet6 fe80::216:3eff:fe23:d5dd/64 scope link  
        valid_lft forever preferred_lft forever  
> 5: veth101.1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue  
>     link/ether 00:16:3e:7e:a5:2e brd ff:ff:ff:ff:ff:ff  
>     inet6 fe80::216:3eff:fe7e:a52e/64 scope link  
        valid_lft forever preferred_lft forever  
>  
> HOST: ip r ls  
>  
> 10.100.225.106 dev veth101.1 scope link  
> 192.168.100.106 dev veth101.0 scope link  
> 192.168.100.0/24 dev eth0 proto kernel scope link src 192.168.100.247  
> 10.100.225.0/24 dev eth1 proto kernel scope link src 10.100.225.247  
> 169.254.0.0/16 dev eth1 scope link  
> default via 192.168.100.1 dev eth0  
>  
> VE: ip a ls
```

>
> 1: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue
> link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
> inet 127.0.0.1/8 scope host lo
> inet6 ::1/128 scope host
> valid_lft forever preferred_lft forever
> 3: venet0: <BROADCAST,POINTOPOINT,NOARP> mtu 1500 qdisc noqueue
> link/void
> 5: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
> link/ether 00:16:3e:69:ce:df brd ff:ff:ff:ff:ff:ff
> inet 192.168.100.106/32 scope global eth0
> inet6 fe80::216:3eff:fe69:cdf/64 scope link
> valid_lft forever preferred_lft forever
> 7: eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
> link/ether 00:16:3e:1d:0d:54 brd ff:ff:ff:ff:ff:ff
> inet 10.100.225.106/32 scope global eth1
> inet6 fe80::216:3eff:fe1d:d54/64 scope link
> valid_lft forever preferred_lft forever
>
> VE: ip r ls
>
> 192.168.100.0/24 dev eth0 scope link
> 10.100.225.0/24 dev eth1 scope link
> default dev eth0 scope link
>
> On 3/14/07, Andrey Mirkin <major@openvz.org> wrote:
>> Hello,
>>
>> Please provide ip configuration (ip a ls) and route tables (ip r ls) for host
>> system and VE.
>>
>> Regards,
>> Andrey
>>
>> On Thursday 15 March 2007 06:16 Julian Yap wrote:
>>> Kind of related to my previous posting.
>>>
>>> I've now setup my 2 VE interfaces using veth method.
>>>
>>> No longer using Source based routing, I've set up my default route to
>>> be 192.168.0.1.
>>>
>>> Pinging on the 192.168.0.x subnet is perfect from the VE.
>>>
>>> Pinging on my VE's 2nd interface (eth1) works OK in burst but at times
>>> just drops packets.
>>>
>>> Running a tcpdump, I see a log of these messages:

> > > 17:10:29.813495 arp who-has machine15 tell 192.168.1.247
> > >
> > > When it does work OK, it's because it's returned this:
> > > 17:14:35.829082 arp who-has machine15 tell 192.168.100.247
> > > 17:14:36.828053 arp who-has machine13 tell machine15
> > > ...
> > > 17:14:36.828082 arp reply machine13 is-at 00:16:3e:7e:a5:2e
> > >
> > > Any clues why it arp's so much?
> > >
> > > - Julian
