
Subject: *SOLVED* VE connect with physical interface
Posted by [gamtech](#) on Mon, 05 Nov 2007 17:58:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dear Community,

I have stuck with one problem and need your help.

We have two HN's with one VE on each. Each HN have 3 interfaces: two merged into bond0 (private network between two HN's) and one eth2 on each node is Internet interface.

Private network is 10.0.0.xxx (HN1 bond0 is 10.0.0.1 and HN2 bond0 is 10.0.0.2).

I would like to transfer data from one VE to another between HN's over private network. Now if I just add 10.0.0.x addresses to each VE (on different HN's), then request goes through private network (bond0), but reply through Internet (eth2). But maybe I mistake and they go completely over eth2.

I've already seen wiki article about venet and veth, but it's production HN's and I don't want completely mess things up doing my research and tests.

Please provide mini how-to add VEs into private network (bond0 interface) without losing Internet access to VEs (I mean primary route through eth2).

Thank you forward!

Subject: Re: VE connect with physical interface
Posted by [khorenko](#) on Tue, 06 Nov 2007 07:52:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

[gamtech](#) wrote on Mon, 05 November 2007 20:58 I would like to transfer data from one VE to another between HN's over private network. Now if I just add 10.0.0.x addresses to each VE (on different HN's), then request goes through private network (bond0), but reply through Internet (eth2). But maybe I mistake and they go completely over eth2.

Let's you have assigned additional private addresses to the VEs:

- VEIP1 for VE1 on HN1
- VEIP2 for VE2 on HN2

Could you please collect the following info?:

- 1) 'ip a l' on both HN's and VEs
- 2) 'ip r l' on both HN's
- 3) 'ip rule l' on both HN's
- 4) 'ip r get \$VEIP1' on the HN2

While pinging VE2 from VE1 run on HN2:

- 5) 'tcpdump -i bond0 -n ip proto \icmp and host \$VEIP1'
- 6) 'tcpdump -i eth2 -n ip proto \icmp and host \$VEIP1'

Thank you,
Konstantin.

Subject: Re: VE connect with physical interface
Posted by [gamtech](#) on Wed, 14 Nov 2007 18:59:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
[root@ve1]# ip a l
1: lo: <LOOPBACK,UP> 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 brd 127.255.255.255 scope host lo
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
    link/void
    inet 127.0.0.1/32 scope host venet0
    inet 143.69.77.8/32 brd 143.69.77.8 scope global venet0:0
    inet 10.10.10.5/32 brd 10.10.10.5 scope global venet0:1
    inet 143.69.77.28/32 brd 143.69.77.28 scope global venet0:10
    inet 143.69.77.29/32 brd 143.69.77.29 scope global venet0:11
    inet 143.69.77.30/32 brd 143.69.77.30 scope global venet0:12
    inet 143.69.77.31/32 brd 143.69.77.31 scope global venet0:13
    inet 143.69.77.32/32 brd 143.69.77.32 scope global venet0:14
    inet 143.69.77.33/32 brd 143.69.77.33 scope global venet0:15
    inet 143.69.77.34/32 brd 143.69.77.34 scope global venet0:16
    inet 143.69.77.35/32 brd 143.69.77.35 scope global venet0:17
    inet 143.69.77.36/32 brd 143.69.77.36 scope global venet0:18
    inet 143.69.77.37/32 brd 143.69.77.37 scope global venet0:19
    inet 143.69.77.20/32 brd 143.69.77.20 scope global venet0:2
    inet 143.69.77.38/32 brd 143.69.77.38 scope global venet0:20
    inet 143.69.77.39/32 brd 143.69.77.39 scope global venet0:21
    inet 143.69.77.40/32 brd 143.69.77.40 scope global venet0:22
    inet 143.69.77.41/32 brd 143.69.77.41 scope global venet0:23
    inet 143.69.77.42/32 brd 143.69.77.42 scope global venet0:24
    inet 143.69.77.43/32 brd 143.69.77.43 scope global venet0:25
    inet 143.69.77.44/32 brd 143.69.77.44 scope global venet0:26
    inet 143.69.77.45/32 brd 143.69.77.45 scope global venet0:27
    inet 143.69.77.46/32 brd 143.69.77.46 scope global venet0:28
    inet 143.69.77.47/32 brd 143.69.77.47 scope global venet0:29
    inet 143.69.77.21/32 brd 143.69.77.21 scope global venet0:3
    inet 143.69.77.48/32 brd 143.69.77.48 scope global venet0:30
    inet 143.69.77.49/32 brd 143.69.77.49 scope global venet0:31
    inet 143.69.77.22/32 brd 143.69.77.22 scope global venet0:4
    inet 143.69.77.23/32 brd 143.69.77.23 scope global venet0:5
    inet 143.69.77.24/32 brd 143.69.77.24 scope global venet0:6
    inet 143.69.77.25/32 brd 143.69.77.25 scope global venet0:7
    inet 143.69.77.26/32 brd 143.69.77.26 scope global venet0:8
    inet 143.69.77.27/32 brd 143.69.77.27 scope global venet0:9
```

```
[root@ve2]# ip a l
1: lo: <LOOPBACK,UP> 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 brd 127.255.255.255 scope host lo
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
```

```
link/void
inet 127.0.0.1/32 scope host venet0
inet 143.69.77.10/32 brd 143.69.77.10 scope global venet0:0
inet 143.69.77.12/32 brd 143.69.77.12 scope global venet0:1
inet 10.10.10.3/32 brd 10.10.10.3 scope global venet0:2
```

```
[root@hn1]# ip a l
```

```
2: lo: <LOOPBACK,UP> 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 brd 127.255.255.255 scope host lo
4: bond0: <BROADCAST,MULTICAST,MASTER,UP> mtu 1500 qdisc noqueue
    link/ether 00:04:23:d6:7e:70 brd ff:ff:ff:ff:ff:ff
    inet 10.10.10.2/24 brd 10.10.10.255 scope global bond0
6: eth0: <BROADCAST,MULTICAST,SLAVE,UP> mtu 1500 qdisc pfifo_fast master bond0 qlen 1000
    link/ether 00:04:23:d6:7e:70 brd ff:ff:ff:ff:ff:ff
8: eth1: <BROADCAST,MULTICAST,SLAVE,UP> mtu 1500 qdisc pfifo_fast master bond0 qlen 1000
    link/ether 00:04:23:d6:7e:70 brd ff:ff:ff:ff:ff:ff
10: eth2: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:04:23:d5:fd:5a brd ff:ff:ff:ff:ff:ff
    inet 143.69.77.6/25 brd 143.69.77.127 scope global eth2
12: eth3: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:04:23:d5:fd:5b brd ff:ff:ff:ff:ff:ff
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
```

```
[root@hn2]# ip a l
```

```
2: lo: <LOOPBACK,UP> 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 brd 127.255.255.255 scope host lo
4: bond0: <BROADCAST,MULTICAST,MASTER,UP> mtu 1500 qdisc noqueue
    link/ether 00:04:23:d5:f1:c0 brd ff:ff:ff:ff:ff:ff
    inet 10.10.10.1/24 brd 192.168.63.255 scope global bond0
6: eth0: <BROADCAST,MULTICAST,SLAVE,UP> mtu 1500 qdisc pfifo_fast master bond0 qlen 1000
    link/ether 00:04:23:d5:f1:c0 brd ff:ff:ff:ff:ff:ff
8: eth1: <BROADCAST,MULTICAST,SLAVE,UP> mtu 1500 qdisc pfifo_fast master bond0 qlen 1000
    link/ether 00:04:23:d5:f1:c0 brd ff:ff:ff:ff:ff:ff
10: eth2: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:04:23:e1:3f:90 brd ff:ff:ff:ff:ff:ff
    inet 143.69.77.4/25 brd 143.69.77.127 scope global eth2
12: eth3: <BROADCAST,MULTICAST> mtu 1500 qdisc noop qlen 1000
    link/ether 00:04:23:e1:3f:91 brd ff:ff:ff:ff:ff:ff
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
    link/void
```

```
[root@hn1]# ip r l
```

```
10.10.10.5 dev venet0 scope link src 10.10.10.2
143.69.77.72 dev venet0 scope link src 10.10.10.2
143.69.77.70 dev venet0 scope link src 10.10.10.2
143.69.77.46 dev venet0 scope link src 10.10.10.2
143.69.77.47 dev venet0 scope link src 10.10.10.2
143.69.77.44 dev venet0 scope link src 10.10.10.2
143.69.77.45 dev venet0 scope link src 10.10.10.2
143.69.77.42 dev venet0 scope link src 10.10.10.2
143.69.77.43 dev venet0 scope link src 10.10.10.2
143.69.77.40 dev venet0 scope link src 10.10.10.2
143.69.77.41 dev venet0 scope link src 10.10.10.2
143.69.77.38 dev venet0 scope link src 10.10.10.2
143.69.77.39 dev venet0 scope link src 10.10.10.2
143.69.77.36 dev venet0 scope link src 10.10.10.2
143.69.77.37 dev venet0 scope link src 10.10.10.2
143.69.77.34 dev venet0 scope link src 10.10.10.2
143.69.77.35 dev venet0 scope link src 10.10.10.2
143.69.77.32 dev venet0 scope link src 10.10.10.2
143.69.77.33 dev venet0 scope link src 10.10.10.2
143.69.77.48 dev venet0 scope link src 10.10.10.2
143.69.77.49 dev venet0 scope link src 10.10.10.2
143.69.77.13 dev venet0 scope link src 10.10.10.2
143.69.77.11 dev venet0 scope link src 10.10.10.2
143.69.77.8 dev venet0 scope link src 10.10.10.2
143.69.77.7 dev venet0 scope link src 10.10.10.2
143.69.77.31 dev venet0 scope link src 10.10.10.2
143.69.77.30 dev venet0 scope link src 10.10.10.2
143.69.77.29 dev venet0 scope link src 10.10.10.2
143.69.77.28 dev venet0 scope link src 10.10.10.2
143.69.77.27 dev venet0 scope link src 10.10.10.2
143.69.77.26 dev venet0 scope link src 10.10.10.2
143.69.77.25 dev venet0 scope link src 10.10.10.2
143.69.77.24 dev venet0 scope link src 10.10.10.2
143.69.77.23 dev venet0 scope link src 10.10.10.2
143.69.77.22 dev venet0 scope link src 10.10.10.2
143.69.77.21 dev venet0 scope link src 10.10.10.2
143.69.77.20 dev venet0 scope link src 10.10.10.2
143.69.77.17 dev venet0 scope link src 10.10.10.2
143.69.77.0/25 dev eth2 proto kernel scope link src 143.69.77.6
10.10.10.0/24 dev bond0 proto kernel scope link src 10.10.10.2
169.254.0.0/16 dev eth2 scope link
default via 143.69.77.1 dev eth2
```

```
[root@hn2]# ip r l
10.10.10.3 dev venet0 scope link src 10.10.10.1
143.69.77.15 dev venet0 scope link src 10.10.10.1
143.69.76.130 dev venet0 scope link src 10.10.10.1
143.69.77.12 dev venet0 scope link src 10.10.10.1
```

```
143.69.77.10 dev venet0 scope link src 10.10.10.1
213.21.225.6 dev venet0 scope link src 10.10.10.1
143.69.77.75 dev venet0 scope link src 10.10.10.1
143.69.77.74 dev venet0 scope link src 10.10.10.1
143.69.77.130 dev venet0 scope link src 10.10.10.1
143.69.77.71 dev venet0 scope link src 10.10.10.1
143.69.77.5 dev venet0 scope link src 10.10.10.1
143.69.77.2 dev venet0 scope link src 10.10.10.1
143.69.77.18 dev venet0 scope link src 10.10.10.1
143.69.77.16 dev venet0 scope link src 10.10.10.1
143.69.77.0/25 dev eth2 proto kernel scope link src 143.69.77.4
10.10.10.0/24 dev bond0 proto kernel scope link src 10.10.10.1
169.254.0.0/16 dev eth2 scope link
default via 143.69.77.1 dev eth2
```

```
[root@hn1]# ip rule l
0:    from all lookup local
32766: from all lookup main
32767: from all lookup default
```

```
[root@hn2]# ip rule l
0:    from all lookup local
32766: from all lookup main
32767: from all lookup default
```

```
[root@hn2]# # ip r get 10.10.10.5
10.10.10.5 dev bond0 src 10.10.10.1
    cache mtu 1500 advmss 1460 hoplimit 64
```

```
[root@hn2]# tcpdump -i bond0 -n ip proto \icmp
20:45:11.169652 IP 143.69.77.8 > 10.10.10.3: icmp 64: echo request seq 0
```

```
[root@hn2]# tcpdump -i eth2 -n ip proto \icmp
20:46:10.271206 IP 10.10.10.3 > 143.69.77.8: icmp 64: echo reply seq 0Request from HN1 to
HN2 goes via bond0, but reply via eth2.
```

Subject: Re: VE connect with physical interface
Posted by [khorenko](#) on Thu, 15 Nov 2007 08:19:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Great, thank you for the info.

The situation is following:

VE1 is pinging VE2 (10.10.10.3).
Routing table on a VE1 says that the packet should go through the venet0 interface (as no any other interfaces exists inside a VE1 :)). As in any other Linux node the source IP for such a

packet is set as a FIRST IP from the interface where this packet is routed.

In your case the first IP is 143.69.77.8 (which can also be seen in tcpdump on a HN2 on bond0 - the packet source IP is 143.69.77.8).

VE2 receives the packet and sends reply to the 143.69.77.8, and HN2 has a route rule that sends such a packet to the eth2. That's it.

So the problem is: VE1 sets "incorrect" (not that as expected) source IP while pinging VE2 (10.10.10.3). This can be handled by adding a route rule in VE1 like:
ip r a 10.10.10.0/24 dev venet0 scope link src 10.10.10.5
You can also take a look at <http://kb.swsoft.com/en/3061>

Hope this helps.

--

Konstantin.

Subject: Re: VE connect with physical interface
Posted by [gamtech](#) on Thu, 15 Nov 2007 08:53:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

Konstantin,
It works! I spasibo.

Subject: Re: VE connect with physical interface
Posted by [khorenko](#) on Thu, 15 Nov 2007 09:13:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

you are welcome.
