

---

Subject: Some VE on one HW node cannot communicate each other (getting Time to live exceeded)

Posted by [Jan Tomasek](#) on Fri, 22 Dec 2006 09:19:42 GMT

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

---

Hello,

I've server with two network cards in following setup:

```
auto eth0
iface eth0 inet static
    address 195.113.233.254
    netmask 255.255.255.0
    network 195.113.233.0
    broadcast 195.113.233.255
    gateway 195.113.233.1
```

```
auto eth0:0
iface eth0:0 inet static
    address 192.168.1.1
    netmask 255.255.255.0
    network 192.168.1.255
    broadcast 192.168.1.255
```

```
auto eth1
iface eth1 inet static
    address 195.113.187.37
    netmask 255.255.255.192
    network 195.113.187.0
    broadcast 195.113.187.63
    post-up ip rule add from 195.113.187.33 table 6
    post-up ip rule add from 195.113.187.35 table 6
    post-up ip route add default dev eth1 via 195.113.187.1 table 6
    post-down ip rule delete from 195.113.187.33 table 6
    post-down ip rule delete from 195.113.187.35 table 6
```

There are three networks on:

- 1) eth0 195.113.233.0/24 primary public network for OpenVZ systems
- 2) eth0:0 192.168.1.0/24 private network for system which should not waste publicaly routable IP but still need to be online
- 3) eth1 195.113.187.0/26 is secondary public network - I was using this range for physical systems which I'm now moving to VE

Everything is working fine, except communication between systems using eth1 and eth0.

Routing tables on HW node:

chlivek:~# ip rule

0: from all lookup 255

32764: from 195.113.187.35 lookup 6

32765: from 195.113.187.33 lookup 6

32766: from all lookup main

32767: from all lookup default

chlivek:~# ip route

195.113.187.33 dev venet0 scope link src 192.168.1.1

192.168.1.2 dev venet0 scope link src 192.168.1.1

195.113.233.253 dev venet0 scope link src 192.168.1.1

195.113.233.252 dev venet0 scope link src 192.168.1.1

195.113.187.0/26 dev eth1 proto kernel scope link src 195.113.187.37

195.113.233.0/24 dev eth0 proto kernel scope link src 195.113.233.254

192.168.1.0/24 dev eth0 proto kernel scope link src 192.168.1.1

default via 195.113.233.1 dev eth0

chlivek:~# route -n

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use
195.113.187.33	0.0.0.0	255.255.255.255	UH	0	0	0
venet0						
192.168.1.2	0.0.0.0	255.255.255.255	UH	0	0	0
venet0						
195.113.233.253	0.0.0.0	255.255.255.255	UH	0	0	0
venet0						
195.113.233.252	0.0.0.0	255.255.255.255	UH	0	0	0
venet0						
195.113.187.0	0.0.0.0	255.255.255.192	U	0	0	0 eth1
195.113.233.0	0.0.0.0	255.255.255.0	U	0	0	0 eth0
192.168.1.0	0.0.0.0	255.255.255.0	U	0	0	0 eth0
0.0.0.0	195.113.233.1	0.0.0.0	UG	0	0	0 eth0

Pinging from system 195.113.187.33 outside:

ermon:~# ping -c 1 195.113.233.254 (HW node)

PING 195.113.233.254 (195.113.233.254) 56(84) bytes of data.

64 bytes from 195.113.233.254: icmp\_seq=1 ttl=64 time=0.035 ms

ermon:~# ping -c 1 195.113.233.1 (Gateway)

PING 195.113.233.1 (195.113.233.1) 56(84) bytes of data.

64 bytes from 195.113.233.1: icmp\_seq=1 ttl=254 time=0.472 ms

ermon:~# ping -c 1 195.113.233.253 (VE using eth0)

PING 195.113.233.253 (195.113.233.253) 56(84) bytes of data.

From 192.168.1.1 icmp\_seq=1 Time to live exceeded

I'm bit curious where that 192.168.1.1 get into path. In FW rules I'm using:

```
iptables -t nat -A POSTROUTING -s 192.168.1.0/24 -o eth0 -j MASQUERADE
echo 1 > /proc/sys/net/ipv4/ip_forward
```

for translating 192.168.1.0/24 onto public 195.113.233.254.

I used [http://wiki.openvz.org/Source\\_based\\_routing](http://wiki.openvz.org/Source_based_routing) as base for setup, but I'm really clueless how to fix it. Please can someone review this setup and suggest me better way how to setup this?

Best regards

--

-----  
Jan Tomasek aka Semik  
<http://www.tomasek.cz/>

---

Subject: Re: Some VE on one HW node cannot communicate each other (getting Time to live exceeded)

Posted by [Jan Tomasek](#) on Thu, 28 Dec 2006 12:24:48 GMT

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

---

Hello again,

I bit moved forward in solving my problem, but still need help.  
Recapitulation first.

I've two network cards, first is connected to 195.113.233.0/24 network:

```
> auto eth0
> iface eth0 inet static
>     address 195.113.233.254
>     netmask 255.255.255.0
>     network 195.113.233.0
>     broadcast 195.113.233.255
>     gateway 195.113.233.1
```

Second is connected to 195.113.187.0/26 network:

```
> iface eth1 inet static
>     address 195.113.187.37
```

```
> netmask 255.255.255.192
> network 195.113.187.0
> broadcast 195.113.187.63
> post-up /etc/network/post-up-eth1
> post-down /etc/network/post-down-eth1
```

Everything is in this setup working, except of communication between VE systems using eth1 and eth0.

Ping from VE (195.113.187.33) using eth1 to system using eth0 (195.113.233.253):

```
ermon:~# ping 195.113.233.253
PING 195.113.233.253 (195.113.233.253) 56(84) bytes of data.
From 195.113.187.37 icmp_seq=1 Time to live exceeded
From 195.113.187.37 icmp_seq=2 Time to live exceeded
```

Pinging to everywhere else is ok.

I discovered that I have to add route for network 195.113.233.0/24 to be routed to eth1 and GW 195.113.187.1 default is to eth0, which sends that TTL exceeded error.

But after setting that route, I'm still not able reach boxes using 195.113.233.0/24 on local eth0 interface. For this I've setup route via venet0 interface.

File "/etc/network/post-up-eth1":

```
#!/bin/bash
```

```
ip rule add from 195.113.187.33 table 6
ip rule add from 195.113.187.35 table 6
ip route add default dev eth1 via 195.113.187.1 table 6
ip route add 195.113.233.0/24 dev eth1 via 195.113.187.1 table 6
ip route add 195.113.233.252 dev venet0 table 6
ip route add 195.113.233.253 dev venet0 table 6
```

```
# Ignore any error of above commands
exit 0;
```

File "/etc/network/post-down-eth1":

```
#!/bin/bash
```

```
ip rule delete from 195.113.187.33 table 6
ip rule delete from 195.113.187.35 table 6
ip route delete 195.113.233.252 dev venet0 table 6
```

```
ip route delete 195.113.233.253 dev venet0 table 6
```

```
# Ignore any error of above commands  
exit 0;
```

```
chlivek:~# ip rule
```

```
0:    from all lookup 255  
32764: from 195.113.187.35 lookup 6  
32765: from 195.113.187.33 lookup 6  
32766: from all lookup main  
32767: from all lookup default
```

```
chlivek:~# ip route
```

```
195.113.187.33 dev venet0 scope link src 195.113.233.254  
192.168.1.2 dev venet0 scope link src 195.113.233.254  
195.113.233.253 dev venet0 scope link src 195.113.187.37  
195.113.233.252 dev venet0 scope link src 195.113.233.254  
195.113.187.0/26 dev eth1 proto kernel scope link src 195.113.187.37  
195.113.233.0/24 dev eth0 proto kernel scope link src 195.113.233.254  
default via 195.113.233.1 dev eth0
```

```
chlivek:~# ip route show table 6
```

```
195.113.233.253 dev venet0 scope link  
195.113.233.252 dev venet0 scope link  
195.113.233.0/24 via 195.113.187.1 dev eth1  
default via 195.113.187.1 dev eth1
```

It seems to work but it is bit complicated. If anyone know simpler way please tell me.

Best regards

--

-----  
Jan Tomasek aka Semik  
<http://www.tomasek.cz/>

---