
Subject: Problems connecting to the outside and back
Posted by [Obi_Wan](#) on Mon, 06 Oct 2008 17:51:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

I don't know how to describe it properly however when I for example try an apt-get update from a container it starts hitting some of the repositories however at some point it just starts hanging at "Waiting for headers". All of the containers have real ips which are routed to my box. Is a problem like that known? Iptables accepts for testing purposes all connections and allows all traffic.

My kernel version:

Linux scs001 2.6.24-6-fza-amd64 #1 SMP Mon May 12 19:35:09 UTC 2008 x86_64 GNU/Linux

Thanks for help.

Edit: I forgot some information. I'm using the venet network module and all my iptables are completely empty. My first guess would be that some of the pakets are only getting routed to the host node and not any further.

Subject: Re: Problems connecting to the outside and back
Posted by [Obi_Wan](#) on Tue, 07 Oct 2008 11:12:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

I've had a look at the tcpdump and my hoster had a look at the router. Somehow my VE with the ip ending 178 sends out packets however the external server it is the host node with ip 176 and sends packets back to 176. This could be the reason why the VE and the external server can't really communicate with each other.

Is a problem like this known? The funny thing is there are connections which actually work. I can connect via port 6667 or 5223 and they work perfectly. Ah and furthermore when I set up the VEs I had to change the netmask for the ips to 255.255.255.224 otherwise communication was impossible.

Would be good if someone could help me

Thanks

Subject: Re: Problems connecting to the outside and back
Posted by [maratrus](#) on Tue, 07 Oct 2008 15:47:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

Do I describe your situation correctly?

- Your VE uses venet0 interface and has IP1=*.*.178.
- You've tried to communicate with the NODE2 which has IP2=*.*.176
- tcpdump shows on the NODE2 that packets are coming from your VE
- but NODE2 sends the packets back to the IP2=*.*.176 instead of IP1 = *.*.178

Am I right?

If yes, this seems to be non OpenVZ problem. Please check iptables/route rules on the Node2

Subject: Re: Problems connecting to the outside and back

Posted by [Obi_Wan](#) on Tue, 07 Oct 2008 16:54:53 GMT

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

My VE uses venet0 with IP *.177 sorry for that
My HN uses *.176 to eth0

VE wants to connect to NODE2. Now Node2 sends packets back to 176 and not 177. I noticed now that the problem happens with many different servers and even clients now.

My IPs have the subnet 255.255.255.224 which is configured inside my eth0. I had to configure the interface venet0 to that subnet as well. However the "route" tells me this:

```
*176 *          255.255.255.255 UH  0  0  0 venet0
```

Could that be the problem?

Edit: I now changed all the subnets back to 255.255.255.255 in the VEs how it was originally. I can connect to the internet now however I still won't get a response. e.g. apt-get update stays at "waiting for headers".

Subject: Re: Problems connecting to the outside and back

Posted by [maratrus](#) on Wed, 08 Oct 2008 07:25:06 GMT

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

Hello,

Quote:

```
*176 * 255.255.255.255 UH 0 0 0 venet0
```

Yes, that's ok. You shouldn't change the netmask inside the VE it should be equal to 32 (255.255.255.255). venet0 is point-to-point connection. The one side is your HN the other side is

VE.

Quote:

Now Node2 sends packets back to 176 and not 177.

And what kind of packets Node2 receives? I mean, what source address that packets have? Do they have *.*.176(your HN) or *.*.177 (your VE) source address? If *.*.176 that's mean that you have nat enabled on your HN please check it carefully.

The exact tcpdump output would be very useful.

Subject: Re: Problems connecting to the outside and back

Posted by [Obi_Wan](#) on Wed, 08 Oct 2008 11:32:13 GMT

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

TCPDUMP

Quote:scs001:/home/klaast# tcpdump -nne host 195.71.9.196

tcpdump: verbose output suppressed, use -v or -vv for full protocol decode

listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes

13:31:12.517977 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 66:
81.2.189.176.53285 > 195.71.9.196.80: F 1746231632:1746231632(0) ack 3540375906 win 71
<nop,nop,timestamp 1624376081 798348372>

13:31:12.575201 00:1b:0d:e4:36:c0 > 00:1e:c9:dc:5a:62, ethertype IPv4 (0x0800), length 66:
195.71.9.196.80 > 81.2.189.176.53285: . ack 1 win 71 <nop,nop,timestamp 798362751
1624376081>

13:31:13.135589 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 74:
81.2.189.176.53305 > 195.71.9.196.80: S 2645258124:2645258124(0) win 5840 <mss
1460,sackOK,timestamp 1624376699 0,nop,wscale 7>

13:31:13.153041 00:1b:0d:e4:36:c0 > 00:1e:c9:dc:5a:62, ethertype IPv4 (0x0800), length 74:
195.71.9.196.80 > 81.2.189.176.53305: S 3590191446:3590191446(0) ack 2645258125 win 5792
<mss 1460,sackOK,timestamp 798362896 1624376699,nop,wscale 7>

13:31:13.153085 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 66:
81.2.189.176.53305 > 195.71.9.196.80: . ack 1 win 46 <nop,nop,timestamp 1624376716
798362896>

13:31:13.153189 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 403:
81.2.189.176.53305 > 195.71.9.196.80: P 1:338(337) ack 1 win 46 <nop,nop,timestamp
1624376716 798362896>

13:31:13.172508 00:1b:0d:e4:36:c0 > 00:1e:c9:dc:5a:62, ethertype IPv4 (0x0800), length 66:
195.71.9.196.80 > 81.2.189.176.53305: . ack 338 win 54 <nop,nop,timestamp 798362900
1624376716>

13:31:13.172562 00:1b:0d:e4:36:c0 > 00:1e:c9:dc:5a:62, ethertype IPv4 (0x0800), length 421:
195.71.9.196.80 > 81.2.189.176.53305: P 1:356(355) ack 338 win 54 <nop,nop,timestamp
798362900 1624376716>

13:31:13.172607 00:1b:0d:e4:36:c0 > 00:1e:c9:dc:5a:62, ethertype IPv4 (0x0800), length 562:
195.71.9.196.80 > 81.2.189.176.53305: P 356:852(496) ack 338 win54 <nop,nop,timestamp
798362900 1624376716>

```
13:31:13.172612 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 66:
81.2.189.176.53305 > 195.71.9.196.80: . ack 356 win 54 <nop,nop,timestamp 1624376736
798362900>
13:31:13.172640 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 66:
81.2.189.176.53305 > 195.71.9.196.80: . ack 852 win 63 <nop,nop,timestamp 1624376736
798362900>
13:31:13.173201 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 242:
81.2.189.176.53305 > 195.71.9.196.80: P 338:514(176) ack 852 win63 <nop,nop,timestamp
1624376736 798362900>
13:31:13.190684 00:1b:0d:e4:36:c0 > 00:1e:c9:dc:5a:62, ethertype IPv4 (0x0800), length 250:
195.71.9.196.80 > 81.2.189.176.53305: P 852:1036(184) ack 514 win 62 <nop,nop,timestamp
798362905 1624376736>
13:31:13.195425 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 259:
81.2.189.176.53305 > 195.71.9.196.80: P 514:707(193) ack 1036 win 71 <nop,nop,timestamp
1624376758 798362905>
13:31:13.415400 00:1e:c9:dc:5a:62 > 00:1b:0d:e4:36:c0, ethertype IPv4 (0x0800), length 259:
81.2.189.176.53305 > 195.71.9.196.80: P 514:707(193) ack 1036 win 71 <nop,nop,timestamp
1624376979 798362905>
13:31:13.432736 00:1b:0d:e4:36:c0 > 00:1e:c9:dc:5a:62, ethertype IPv4 (0x0800), length 78:
195.71.9.196.80 > 81.2.189.176.53305: . ack 707 win 71 <nop,nop,timestamp 798362966
1624376979,nop,nop,sack 1 {514:707}>
```

VE:

```
Quote:oot@scs002:/# apt-get update
Get:1 http://ftp2.de.debian.org etch Release.gpg [386B]
Get:2 http://ftp2.de.debian.org etch/updates Release.gpg [189B]
Ign http://dotdeb.netmirror.org stable Release.gpg
Hit http://ftp2.de.debian.org etch Release
Ign http://dotdeb.netmirror.org stable Release
Ign http://dotdeb.netmirror.org stable/all Packages/DiffIndex
Ign http://dotdeb.netmirror.org stable/all Sources/DiffIndex
Ign http://dotdeb.netmirror.org stable/all Packages
Ign http://dotdeb.netmirror.org stable/all Sources
Hit http://dotdeb.netmirror.org stable/all Packages
Hit http://dotdeb.netmirror.org stable/all Sources
99% [Waiting for headers]
```

Subject: Re: Problems connecting to the outside and back

Posted by [maratrus](#) on Wed, 08 Oct 2008 11:55:09 GMT

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

Ok,

from tcpdump output we can conclude that the packets from 81.2.189.176 are coming but not from your VE.

Please, check your iptables rules:

```
# iptables -t nat -L
```

```
# iptables -t mangle -L
# iptables -t filter -L
(from HN and from inside the VE)
You also can switch off iptables temporarily
```

Subject: Re: Problems connecting to the outside and back
Posted by [Obi_Wan](#) on Wed, 08 Oct 2008 14:43:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

From the HN:

Quote:Chain POSTROUTING (policy ACCEPT)

target	prot	opt	source	destination	
SNAT	0	--	anywhere	anywhere	to:81.2.189.176
SNAT	0	--	anywhere	anywhere	to:81.2.189.176
SNAT	0	--	scs003.st-city.net	anywhere	to:81.2.189.178

Chain OUTPUT (policy ACCEPT)

target	prot	opt	source	destination
scs001:/home/klaast#	iptables	-t	mangle	-L

Chain PREROUTING (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------

Chain INPUT (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------

Chain FORWARD (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------

Chain OUTPUT (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------

Chain POSTROUTING (policy ACCEPT)

target	prot	opt	source	destination
scs001:/home/klaast#	iptables	-t	filter	-L

Chain INPUT (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------

Chain FORWARD (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------

Chain OUTPUT (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------

No Iptables on the VE enabled.

Hm... I can't remember to have enabled these rules. Could these be the reason why my packets

go missing?

Subject: Re: Problems connecting to the outside and back

Posted by [Obi_Wan](#) on Wed, 08 Oct 2008 17:11:51 GMT

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

Hi,

I just disabled these NAT rules and now the right IPs are communicating.

```
Quote:19:09:12.527168 IP 81.2.189.177.53999 > 195.71.9.196.80: F 2389437050:2389437050(0)
ack 3517960015 win 71 <nop,nop,timestamp 1644656090 803425662>
19:09:12.587090 IP 195.71.9.196.80 > 81.2.189.177.53999: . ack 1 win 71 <nop,nop,timestamp
803432602 1644656090>
19:09:13.066110 IP 81.2.189.177.54011 > 195.71.9.196.80: S 2837860495:2837860495(0) win
5840 <mss 1460,sackOK,timestamp 1644656629 0,nop,wscale 7>
19:09:13.083733 IP 195.71.9.196.80 > 81.2.189.177.54011: S 3535821303:3535821303(0) ack
2837860496 win 5792 <mss 1460,sackOK,timestamp 803432727 1644656629,nop,wscale 7>
19:09:13.083824 IP 81.2.189.177.54011 > 195.71.9.196.80: . ack 1 win 46 <nop,nop,timestamp
1644656647 803432727>
19:09:13.083908 IP 81.2.189.177.54011 > 195.71.9.196.80: P 1:338(337) ack 1 win 46
<nop,nop,timestamp 1644656647 803432727>
19:09:13.101084 IP 195.71.9.196.80 > 81.2.189.177.54011: . ack 338 win 54 <nop,nop,timestamp
803432731 1644656647>
19:09:13.101176 IP 195.71.9.196.80 > 81.2.189.177.54011: P 1:356(355) ack 338 win 54
<nop,nop,timestamp 803432731 1644656647>
19:09:13.101291 IP 81.2.189.177.54011 > 195.71.9.196.80: . ack 356 win 54 <nop,nop,timestamp
1644656665 803432731>
19:09:13.101794 IP 195.71.9.196.80 > 81.2.189.177.54011: P 356:852(496) ack 338 win 54
<nop,nop,timestamp 803432731 1644656647>
19:09:13.101826 IP 81.2.189.177.54011 > 195.71.9.196.80: . ack 852 win 63 <nop,nop,timestamp
1644656665 803432731>
19:09:13.102058 IP 81.2.189.177.54011 > 195.71.9.196.80: P 338:514(176) ack 852 win 63
<nop,nop,timestamp 1644656665 803432731>
19:09:13.119898 IP 195.71.9.196.80 > 81.2.189.177.54011: P 852:1036(184) ack 514 win 62
<nop,nop,timestamp 803432735 1644656665>
19:09:13.125095 IP 81.2.189.177.54011 > 195.71.9.196.80: P 514:707(193) ack 1036 win 71
<nop,nop,timestamp 1644656688 803432735>
19:09:13.351031 IP 81.2.189.177.54011 > 195.71.9.196.80: P 514:707(193) ack 1036 win 71
<nop,nop,timestamp 1644656914 803432735>
19:09:13.368264 IP 195.71.9.196.80 > 81.2.189.177.54011: . ack 707 win 71 <nop,nop,timestamp
803432798 1644656914,nop,nop,sack 1 {514:707}>
```

However it is still on "waiting for headers".

Here is my sysctl.conf

```
Quote:kernel.msgmnb = 65536
kernel.msgmax = 65536
kernel.shmmax = 4294967295
kernel.shmall = 268435456
net.ipv4.ip_forward = 1
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.default.accept_source_route = 0
kernel.sysrq = 1
kernel.core_uses_pid = 1
net.ipv4.tcp_syncookies = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
```

Subject: Re: Problems connecting to the outside and back
Posted by [maratrus](#) on Thu, 09 Oct 2008 06:50:53 GMT

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

Hi,

first of all, thanks for the patient.

some more questions please:

- are you able to ping 195.71.9.196 from inside the VE?
- please run tcpdump (during the attempt to communicate with 195.71.9.196) on incoming interface on HN (eth0 for example), venet0 interface on HN and venet0 interface inside your VE simultaneously. We have to make sure that all packets receive VE. I guess packets do receive your VE.
- if everything is ok, please check your repository settings on the node 195.71.9.196. May be the problem is on that node.

Subject: Re: Problems connecting to the outside and back
Posted by [Obi_Wan](#) on Thu, 09 Oct 2008 08:02:47 GMT

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

I haven't got the possibility to get the real output currently but I will try this afternoon the latest tomorrow.

However yesterday I tested tcpdump on the HN and tcpdump on the VE. All packets which arrive the HN arrive the VE. The IP belongs to ftp2.de.debian.org and I have the same result with

ftp.de.debian.org and ftp.us.debian.org

After some time the connection just times out. Now I noticed I have many errors on the eth0 interface on the receiving section.

Sorry for not being able to give you exact information currently. I will try my best but maybe this will give you some further ideas.

Edit: Ah and a similar result with other connections. When I connect an ircd on my VE with an external ircd I'm unable to get a full response so they are unable to sync.

Subject: Re: Problems connecting to the outside and back

Posted by [Obi_Wan](#) on Fri, 10 Oct 2008 21:42:57 GMT

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

Here my new information. The ping from inside the VE works. I checked the tcpdump and everything seems to arrive.

The other nodes repository should be ok because it is ftp2.de.debian.org one of the official mirrors. As I said ftp.de.debian.org and ftp.us.debian.org the same. It will stay with "waiting for headers" until the connection times out. However now I found a very quick display that something was actually transferred at some bytes/s and then dies. I have also noticed that everytime I do some apt-get update or something which doesn't work I receive many RX errors on eth0. My hoster says everything is correctly connectet :-/

Thanks for your help so far.

Edit: I investigated a bit further and I found out that these RX errors I get are all "frame too long" errors.
