
Subject: no connectivity from inside VE after migration

Posted by [goeldi](#) on Mon, 14 Jan 2008 10:22:55 GMT

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

OK, this really seems to haunt me

I have 2 hardware nodes on different networks. Node A and Node B. On Node A are 4 VEs running Centos 4 and 5:

vzlist

VEID	NPROC	STATUS	IP_ADDR	HOSTNAME
1233	49	running	79.47.111.233	mail3.bla.net
3226	43	running	79.47.193.226	bla.ch
3229	28	running	79.47.193.229	mail.bla.ch
3230	4	running	79.47.193.230	plone3.bla.net

VE 1233 with Centos 4

VE 3226 with Centos 5

VE 3229 with Centos 4

VZ 3230 with Centos 5

This Node (A) had a disk error (software RAID-1) so I migrated VE 3226 to HN B. The other ones were vdzipped and copied to a backup server.

VE 3226 became another IP address on HN B (because HN B is on a completely different network). Everything was OK.

After the disk change I reinstalled VZ on HN A, going to Centos 5 and LVM.

Now I can do everything with VE 1233: entering and pinging to ip and domains.

VE 3226, 3229 and 3230 have no connectivity outside of the HN: pinging to ip or domain does not work. I stopped iptables on VE and on HN. Does not change anything.

Here I paste the data from HN and VE 3226:

HN

ip a l

```
2: lo: <LOOPBACK,UP,LOWER_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 scope host lo
    inet6 ::1/128 scope host
```

```
    valid_lft forever preferred_lft forever
```

```
4: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:e0:4c:10:05:18 brd ff:ff:ff:ff:ff:ff
    inet 213.239.194.78/27 brd 213.239.194.95 scope global eth0
    inet6 fe80::2e0:4cff:fe10:518/64 scope link
    valid_lft forever preferred_lft forever
```

```
6: sit0: <NOARP> mtu 1480 qdisc noop
  link/sit 0.0.0.0 brd 0.0.0.0
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
  link/void
```

```
# ip r l
79.47.193.229 dev venet0 scope link
79.47.193.230 dev venet0 scope link
79.47.193.226 dev venet0 scope link
79.47.111.233 dev venet0 scope link
213.239.194.64/27 dev eth0 proto kernel scope link src 213.239.194.78
169.254.0.0/16 dev eth0 scope link
default via 213.239.194.65 dev eth0
```

```
# iptables -L
Chain INPUT (policy ACCEPT)
target    prot opt source                destination
RH-Firewall-1-INPUT  all  --  anywhere              anywhere
```

```
Chain FORWARD (policy ACCEPT)
target    prot opt source                destination
```

```
Chain OUTPUT (policy ACCEPT)
target    prot opt source                destination
```

```
Chain RH-Firewall-1-INPUT (1 references)
target    prot opt source                destination
ACCEPT    all  --  anywhere              anywhere
ACCEPT    icmp --  anywhere              anywhere      icmp any
ACCEPT    esp  --  anywhere              anywhere
ACCEPT    ah   --  anywhere              anywhere
ACCEPT    udp  --  anywhere              224.0.0.251      udp dpt:mdns
ACCEPT    udp  --  anywhere              anywhere          udp dpt:ipp
ACCEPT    all  --  anywhere              anywhere          state RELATED,ESTAB
                                LISHED
ACCEPT    tcp  --  anywhere              anywhere          state NEW tcp dpt:n
                                dmp
ACCEPT    tcp  --  anywhere              anywhere          state NEW tcp dpt:s
                                sh
ACCEPT    tcp  --  212-98-47-70.static.adslpremium.ch anywhere          stat
                                e NEW tcp dpt:19150
ACCEPT    tcp  --  192.168.2.3            anywhere          state NEW tcp dpt:1
                                9150
REJECT    all  --  anywhere              anywhere          reject-with icmp-ho
                                st-prohibited
```

```
# iptables -t nat -L
Chain PREROUTING (policy ACCEPT)
```

target prot opt source destination

Chain POSTROUTING (policy ACCEPT)

target prot opt source destination

Chain OUTPUT (policy ACCEPT)

target prot opt source destination

No file /proc/sys/net/ipv4/ip_forwarding on hardware node!

Inside VE:

ip a l

1: lo: <LOOPBACK,UP,LOWER_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 scope host lo

inet6 ::1/128 scope host

valid_lft forever preferred_lft forever

3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue

link/void

inet 127.0.0.1/32 scope host venet0

inet 79.47.193.226/32 brd 79.47.193.226 scope global venet0:0

ip r l

192.0.2.0/24 dev venet0 scope host

169.254.0.0/16 dev venet0 scope link

default via 192.0.2.1 dev venet0

iptables -L

Chain INPUT (policy ACCEPT)

target prot opt source destination

ACCEPT tcp -- anywhere anywhere

tcp flags:ACK/ACK

ACCEPT all -- anywhere anywhere

state ESTABLISHED

ACCEPT all -- anywhere anywhere

state RELATED

ACCEPT udp -- anywhere anywhere

udp spt:domain dpts:1024:65535

ACCEPT icmp -- anywhere anywhere

icmp echo-reply

ACCEPT icmp -- anywhere anywhere

icmp destination-unreachable

ACCEPT icmp -- anywhere anywhere

icmp source-quench

ACCEPT icmp -- anywhere anywhere

icmp time-exceeded

ACCEPT icmp -- anywhere anywhere

icmp parameter-problem

ACCEPT tcp -- anywhere anywhere

tcp dpt:ssh

ACCEPT tcp -- anywhere anywhere

tcp dpt:smtp

ACCEPT tcp -- anywhere anywhere

tcp dpt:http

ACCEPT tcp -- anywhere anywhere

tcp dpt:https

ACCEPT tcp -- anywhere anywhere

tcp dpt:pop3

ACCEPT tcp -- anywhere anywhere

tcp dpt:imap

```
ACCEPT  tcp -- anywhere      anywhere      tcp dpts:ndmp:10010
ACCEPT  tcp -- anywhere      anywhere      tcp dpt:dnp state NEW
REJECT  all -- anywhere      anywhere      reject-with icmp-port-unreachable
```

Chain FORWARD (policy ACCEPT)

```
target  prot opt source      destination
```

Chain OUTPUT (policy ACCEPT)

```
target  prot opt source      destination
```

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

Chain PREROUTING (policy ACCEPT)

```
target  prot opt source      destination
```

Chain POSTROUTING (policy ACCEPT)

```
target  prot opt source      destination
```

Chain OUTPUT (policy ACCEPT)

```
target  prot opt source      destination
```

Also on VE no /proc/sys/net/ipv4/ip_forwarding exists.
