
Subject: VE loses connectivity

Posted by [mow](#) on Sun, 15 Mar 2009 03:54:01 GMT

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

Hi All,

I have one ethernet device, 16 ip addresses in one subnet (call it subnet1) , 4 ip addresses in a different subnet (subnet2). My HN ip address is on subnet1. I have no problem about assigning or connecting to VE's with different subnets (in /etc/vz/vz.conf NEIGHBOUR_DEVS=all) except, in a period of time some of my VE's (both on subnet1 and subnet2) loses connections.

To define the problem, while VE is not responding I did a tcpdump for VE's ip address (77.xxx.xxx.170) on HN and make a traceroute from outside of my network, results below;

```
[root@77-xxx-xxx-162 ~]# tcpdump -r dump70
reading from file dump70, link-type EN10MB (Ethernet)
05:02:04.469826 arp who-has 77-xxx-xxx-170.myhostname.net (Broadcast) tell
host.mygateway.com
05:02:05.085170 arp reply 77-xxx-xxx-170.myhostname.net is-at 00:1f:d0:85:e0:93 (oui Unknown)
05:02:09.464329 arp who-has 77-xxx-xxx-170.myhostname.net (Broadcast) tell
host.mygateway.com
05:02:09.601086 arp reply 77-xxx-xxx-170.myhostname.net is-at 00:1f:d0:85:e0:93 (oui Unknown)
05:02:14.464240 arp who-has 77-xxx-xxx-170.myhostname.net (Broadcast) tell
host.mygateway.com
05:02:14.992004 arp reply 77-xxx-xxx-170.myhostname.net is-at 00:1f:d0:85:e0:93 (oui Unknown)
```

Why there is so much arp queries made by gateway?

As my knowledge it must be only one time (who-has <-> arp_reply only 1 time) isnt it?

When connection goes, If I manually type; (on HN)

```
#arp send -U -i 77.xxx.xxx.170 -c 1 eth0
```

connection comes back. Do you have any ideas? Can this happen because of a misconfigured switch/router (may be like saving/using a static arp table)?

Other details below;

```
[root@77-xxx-xxx-162 ~]# uname -a
```

```
Linux 77-xxx-xxx-162.myhostname.net 2.6.18-92.1.18.el5.028stab060.2 #1 SMP Tue Jan 13
11:38:36 MSK 2009 x86_64 x86_64 x86_64 GNU/Linux
```

```
[root@77-xxx-xxx-162 ~]# ip r l
```

```
77.xxx.xxx.170 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.144.69 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.144.68 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.171 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.168 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.169 dev venet0 scope link src 77.xxx.xxx.162
```

```
77.xxx.xxx.174 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.175 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.172 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.173 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.144.76 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.163 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.176 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.177 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.166 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.167 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.144.75 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.164 dev venet0 scope link src 77.xxx.xxx.162
77.xxx.xxx.160/27 dev eth0 proto kernel scope link src 77.xxx.xxx.162
169.254.0.0/16 dev eth0 scope link
default via 77.xxx.xxx.161 dev eth0
```

sysctl variables;

```
net.ipv4.ip_forward = 1
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.default.accept_source_route = 0
kernel.sysrq = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.all.log_martians = 1
net.ipv4.icmp_echo_ignore_broadcasts = 1
net.ipv4.icmp_echo_ignore_all = 1
kernel.core_uses_pid = 1
net.ipv4.tcp_syncookies = 1
net.ipv4.tcp_max_syn_backlog = 2048
net.ipv4.tcp_synack_retries = 2
net.ipv4.tcp_fin_timeout = 15
net.ipv4.tcp_keepalive_time = 1800
net.ipv4.tcp_window_scaling = 0
net.ipv4.tcp_sack = 0
net.ipv4.tcp_timestamps = 0
net.ipv4.icmp_ignore_bogus_error_responses = 1
```

```
[root@77-xxx-xxx-162 ~]# cat /etc/vz/vz.conf
```

```
## Global parameters
```

```
VIRTUOZZO=yes
```

```
LOCKDIR=/vz/lock
```

```
DUMPDIR=/vz/dump
```

```
VE0CPUUNITS=1000
```

```
## Logging parameters
```

```
LOGGING=yes
LOGFILE=/var/log/vzctl.log
LOG_LEVEL=0
VERBOSE=0
```

```
## Disk quota parameters
DISK_QUOTA=yes
VZFASTBOOT=no
```

```
# Disable module loading. If set, vz initscript do not load any modules.
#MODULES_DISABLED=yes
```

```
# The name of the device whose IP address will be used as source IP for CT.
# By default automatically assigned.
VE_ROUTE_SRC_DEV="eth0"
```

```
# Controls which interfaces to send ARP requests and modify APR tables on.
NEIGHBOUR_DEVS=all
```

```
## Template parameters
TEMPLATE=/vz/template
```

```
## Defaults for containers
VE_ROOT=/vz/root/$VEID
VE_PRIVATE=/vz/private/$VEID
CONFIGFILE="vps.basic"
DEF_OSTEMPLATE="fedora-core-4"
```

```
## Load vzwdog module
VZWDOG="no"
```

```
## IPv4 iptables kernel modules
IPTABLES="ipt_REJECT ipt_tos ipt_limit ipt_multiport iptable_filter iptable_mangle ipt_TCPMSS
ipt_tcpmss ipt_ttl ipt_length iptable_nat"
```

```
## Enable IPv6
IPV6="no"
```

```
## IPv6 ip6tables kernel modules
IP6TABLES="ip6_tables ip6table_filter ip6table_mangle ip6t_REJECT"
```

Subject: Re: VE loses connectivity
Posted by [maratrus](#) on Tue, 17 Mar 2009 09:41:12 GMT

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

Hi,

Quote:

As my knowledge it must be only one time (who-has <-> arp_reply only 1 time) isnt it?

Generally speaking - yes.

But it might have happened that arp-replies didn't reach host.mygateway.com.

By the way, does tcpdump output was taken from eth0 interface?

Does 00:1f:d0:85:e0:93 belong to eth0?

Do you have a chance to check if arp-reply actually reach host.mygateway.com?

Subject: Re: VE loses connectivity

Posted by [mow](#) on Tue, 17 Mar 2009 09:50:05 GMT

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

Yes, dumped from HN-eth0 and mac (00:1f:d0:85:e0:93) belongs to eth0.

I cant control the gateway and I dont know how to check this, but when I manually broadcast arp data every 5 minutes it is working.

Subject: Re: VE loses connectivity

Posted by [maratrus](#) on Wed, 18 Mar 2009 08:54:49 GMT

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

Hello,

it doesn't look like an OpenVZ-specific problem but who knows ...

The only difference between

Quote:

05:02:05.085170 arp reply 77-xxx-xxx-170.myhostname.net is-at 00:1f:d0:85:e0:93 (oui Unknown)

and

Quote:

arp send -U -i 77.xxx.xxx.170 -c 1 eth0

is that the former network packet is sent with particular destination MAC address (BTW, check it with "-e" option of tcpdump utility) and the latter with broadcast MAC.
