
Subject: Network down

Posted by Drago **on Mon, 15 Mar 2010 13:24:42 GMT**

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

Hello.

The problem is, that as containers are working, the network to someone or more than one stops. It is not necessary that the container is one and the same everytime. When I run ping to the container from the host node, there is no reply. I can enter the container with "vzctl enter XXX", but the problem stays. The problem is fixed when I execute "/sbin/ifdown venet0 && /sbin/ifup venet0".

Sometimes this doesn't help, because in 1 min, another container could stop. Sometimes it works normally for day or two without any problems, but after that it could start happening every 5 mins.

The kernel is Linux ufo.myhost.com 2.6.18-164.11.1.el5.028stab068.3 #1 SMP Wed Feb 17

15:22:30 MSK 2010 x86_64 x86_64 x86_64 GNU/Linux

iptables :

mangle and nat are empty and with -P ACCEPT
in filter have only in FORWARD -J ACCEPT rules
INPUT and OUTPUT are empty too

Subject: Re: Network down

Posted by maratrus **on Fri, 19 Mar 2010 13:24:21 GMT**

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

Hello,

if I'm not mistaken you wrote to OpenVZ users mailing list about the same problem. I answered it several days ago so I was hoping that you have to receive an answer. Anyway, here was the answer

Quote:

Hi,

as far as I understand, your network configuration is based on simple venet0 interface.

Is that true? I suppose that you are faced with arp-problem but could you please elaborate

your network configuration a little bit so one can understand what the exact environment is.

It may be important if you are using several route tables.

"ip a l", "ip route list table all", "ip rule list", "arp -n" would be enough I suppose.

Let me give you a hint so that you will be able to cope with the problem by yourself.

venet0 is working according the following principle. If a remote machine is willing to communicate

with a VE it send "arp-who has" request. This type of request reaches a HN and the HN is sending

"arp reply" to the remote machine (that's why "arp -n" output should contain information about VE).

Then the remote machine sends network packets to the HN but because of the additional route

(see "ip route list" output) all packets are going inside VE through the HN. That's the principle of venet0

interface.

To catch the problem I recommend you using "tcpdump" utility.

Subject: Re: Network down

Posted by [Drago](#) on Fri, 19 Mar 2010 15:44:47 GMT

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

Hi,

I use venet0 interface and here is output result of "ip a l", "ip route list table all", "ip rule list", "arp -n"

I sniff with tcpdump but this is hard because I don't know when the container will be down.

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 cbq qlen 100
    link/ether 00:30:48:ca:c5:ac brd ff:ff:ff:ff:ff:ff
    inet XXX.XX.247.2/30 brd XXX.XX.247.3 scope global eth0
        inet6 fe80::230:48ff:fea:c5ac/64 scope link
            valid_lft forever preferred_lft forever
6: eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop qlen 1000
    link/ether 00:30:48:ca:c5:ad brd ff:ff:ff:ff:ff:ff
1: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc cbq
    link/void
```

```
ip route list table all
```

```
85.14.28.114 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.12 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.44 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.13 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.10 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.43 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.11 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.42 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.41 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.40 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.9 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.39 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.6 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.23 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.38 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.4 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.5 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.36 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.35 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.34 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.96 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.33 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.16 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.32 dev venet0 scope link src XXX.XX.247.2
XXX.XX.247.0/30 dev eth0 proto kernel scope link src XXX.XX.247.2
169.254.0.0/16 dev eth0 scope link
default via XXX.XX.247.1 dev eth0
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
local XXX.XX.247.2 dev eth0 table 255 proto kernel scope host src XXX.XX.247.2
broadcast XXX.XX.247.3 dev eth0 table 255 proto kernel scope link src XXX.XX.247.2
broadcast XXX.XX.247.0 dev eth0 table 255 proto kernel scope link src XXX.XX.247.2
broadcast 127.0.0.0 dev lo table 255 proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table 255 proto kernel scope host src 127.0.0.1
fe80::/64 dev eth0 metric 256 expires 21152872sec mtu 1500 advmss 1440 hoplimit
4294967295
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
local ::1 via :: dev lo table 255 proto none metric 0 mtu 16436 advmss 16376 hoplimit
4294967295
local fe80:: via :: dev lo table 255 proto none metric 0 mtu 16436 advmss 16376 hoplimit
4294967295
```

```
local fe80::230:48ff:feca:c5ac via :: dev lo table 255 proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
ff00::/8 dev eth0 table 255 metric 256 expires 21152872sec mtu 1500 advmss 1440 hoplimit 4294967295
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
```

ip rule list

```
0: from all lookup 255
32766: from all lookup main
32767: from all lookup default
```

arp -n

Address	HWtype	HWaddress	Flags Mask	Iface
XXX.XX.247.1	ether	00:1E:13:E4:79:C0	C	eth0
XXX.XX.247.36	*	*	MP	eth0
XXX.XX.247.6	*	*	MP	eth0
XXX.XX.247.23	*	*	MP	eth0
XXX.XX.247.96	*	*	MP	eth0
XXX.XX.247.5	*	*	MP	eth0
XXX.XX.247.39	*	*	MP	eth0
XXX.XX.247.38	*	*	MP	eth0
XXX.XX.247.4	*	*	MP	eth0
XXX.XX.247.33	*	*	MP	eth0
XXX.XX.247.32	*	*	MP	eth0
85.14.28.114	*	*	MP	eth0
XXX.XX.247.35	*	*	MP	eth0
XXX.XX.247.16	*	*	MP	eth0
XXX.XX.247.34	*	*	MP	eth0
XXX.XX.247.44	*	*	MP	eth0
XXX.XX.247.13	*	*	MP	eth0
XXX.XX.247.12	*	*	MP	eth0
XXX.XX.247.11	*	*	MP	eth0
XXX.XX.247.41	*	*	MP	eth0
XXX.XX.247.10	*	*	MP	eth0
XXX.XX.247.40	*	*	MP	eth0
XXX.XX.247.43	*	*	MP	eth0
XXX.XX.247.9	*	*	MP	eth0
XXX.XX.247.42	*	*	MP	eth0

and here is log with info for hosts down :

```
Mon Mar 15 13:21:53 2010: Mon Mar 15 13:22:41 2010: Mon Mar 15 13:23:13 2010:
```

XXX.XX.247.36
Mon Mar 15 13:23:18 2010: XXX.XX.247.36
Mon Mar 15 13:23:28 2010: XXX.XX.247.36
Mon Mar 15 13:26:19 2010: Down Hosts : XXX.XX.247.36
Mon Mar 15 13:26:42 2010: Down Hosts :
XXX.XX.247.5;XXX.XX.247.36;XXX.XX.247.35;XXX.XX.247.39;XXX.X
X.247.11;XXX.XX.247.16;XXX.XX.247.6;XXX.XX.247.4;XXX.XX.247.
38;XXX.XX.247.34;XXX.XX.247.40;XXX.XX.247.32
Mon Mar 15 13:27:49 2010: Down Hosts : XXX.XX.247.36
Mon Mar 15 13:33:02 2010: Down Hosts : XXX.XX.247.36
Mon Mar 15 13:34:02 2010: Down Hosts : XXX.XX.247.36
Mon Mar 15 13:46:04 2010: Down Hosts : XXX.XX.247.34 ; XXX.XX.247.23 ; XXX.XX.247.42
Mon Mar 15 14:08:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Mon Mar 15 14:40:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Mon Mar 15 14:54:03 2010: Down Hosts : XXX.XX.247.36
Mon Mar 15 15:01:02 2010: Down Hosts : XXX.XX.247.36
Mon Mar 15 15:32:02 2010: Down Hosts : XXX.XX.247.36
Mon Mar 15 16:06:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Mon Mar 15 22:38:02 2010: Down Hosts : XXX.XX.247.35
Tue Mar 16 00:22:06 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:23:18 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:24:23 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:25:36 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:26:12 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:27:24 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:28:18 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:29:17 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:30:22 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:31:11 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:32:16 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 00:33:09 2010: Down Hosts : XXX.XX.247.11
Tue Mar 16 01:38:02 2010: Down Hosts : XXX.XX.247.16
Tue Mar 16 04:54:02 2010: Down Hosts : XXX.XX.247.6
Tue Mar 16 04:55:08 2010: Down Hosts : XXX.XX.247.6
Tue Mar 16 04:56:04 2010: Down Hosts : XXX.XX.247.6
Wed Mar 17 15:31:02 2010: Down Hosts : XXX.XX.247.36
Wed Mar 17 15:32:02 2010: Down Hosts : XXX.XX.247.36
Wed Mar 17 17:26:02 2010: Down Hosts : XXX.XX.247.38
Wed Mar 17 22:03:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 01:27:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:28:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:29:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:30:04 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:31:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96

Thu Mar 18 01:32:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:33:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:34:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:35:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:39:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:40:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:41:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:42:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:44:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 01:45:03 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.96
Thu Mar 18 03:01:03 2010: Down Hosts : XXX.XX.247.34 ; XXX.XX.247.23
Thu Mar 18 03:20:07 2010: Down Hosts : XXX.XX.247.35 ; XXX.XX.247.39 ; XXX.XX.247.11 ;
XXX.XX.247.16 ; XXX.XX.247.6 ; XXX.XX.247.96
Thu Mar 18 06:01:03 2010: Down Hosts : XXX.XX.247.34 ; XXX.XX.247.23
Thu Mar 18 08:01:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 08:32:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 10:58:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 10:59:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 11:00:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 11:01:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 11:02:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 11:03:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6
Thu Mar 18 11:04:05 2010: Down Hosts : XXX.XX.247.39 ; XXX.XX.247.11 ; XXX.XX.247.16 ;
XXX.XX.247.6

Thanks.

Subject: Re: Network down
Posted by [maratus](#) on Fri, 19 Mar 2010 16:39:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

Quote:

I sniff with tcpdump but this is hard because I don't know when the container will be down.

at the moment everything is fine and I can't see anything strange in provided output. So you'd

better catch a moment when a particular VE is down and then:

- compare the provided output with output that will be at that moment (or just post it here)
- run tcpdump on eth0 and venet0 interfaces on the HN as well as on the venet0 interface inside the VE
- check sysctls http://wiki.openvz.org/Quick_installation#sysctl
- check routes/iptables inside the VE

P.S. BTW, who writes "Down Hosts"? What kind of program puts that messages in logs and how it indicates that hosts are really down?

Subject: Re: Network down

Posted by [Drago](#) on Fri, 19 Mar 2010 17:15:53 GMT

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

cat /etc/sysctl.conf

```
net.ipv4.ip_forward = 1
net.ipv6.conf.default.forwarding = 1
net.ipv6.conf.all.forwarding = 1
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.rp_filter = 1
kernel.sysrq = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
net.ipv4.tcp_syncookies = 1
net.ipv4.tcp_max_syn_backlog = 2048
net.ipv4.tcp_synack_retries = 2
net.ipv4.tcp_abort_on_overflow = 1
net.ipv4.icmp_echo_ignore_broadcasts=1
net.ipv4.conf.default.forwarding=1
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.send_redirects = 0
```

```
net.ipv4.conf.all.arp_ignore=1
net.ipv4.conf.all.arp_announce=2
```

and script who write host is down

simple perl script

```
#!/usr/bin/perl
#
use vars qw/ $p, @p, @inactive, @skip, @ping, %p /;
```

```

use Net::Ping;

$p =~ s/\s+/ /;
$p =~ s/\s+$/g;
@p = split "\n", $p;
for ( @p ) {
    $_ =~ s/\s\C+$/g;
}

$p = `/usr/sbin/vzlist -H -o ip -S`;
$p =~ s/\s+/ /;
$p =~ s/\s+$/g;
@skip = split "\n", $p;
for ( @skip ) {
    $_ =~ s/\s\C+$/g;
    $p{$_} = 1;
}

for ( @p ) {
    if ($p{$_}) {
    } else {
        push @ping, $_;
    }
}

$p = Net::Ping->new("icmp");
foreach $host (@ping)
{
    push @inactive, $host unless $p->ping($host, 1);
}
$p->close();

if ( @inactive ) {
    open (FILE,'>>/var/log/ping_hosts.log');
    print FILE localtime().": Down Hosts :".(join ' ; ', @inactive)."\n";
    close FILE;
    $p = `/sbin/ifdown venet0`;
    sleep(2);
    $p = `/sbin/ifup venet0`;
}

```

when scripts write down hosts are realy down /i cant ping them from HN with ping xxx.xx.xxx.x/.

and inside VE no any iptable rules. And routes are clean.

