
Subject: Fully access to NIC

Posted by [TSchilli](#) on Mon, 31 Jul 2006 09:30:28 GMT

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

Hello,

we try to use ARP-Watch and similar tools in a VPS, but currently it doesn't work.

Is there any possibility to grant those tools fully access to an interface?

I've read about moving the interface into the VPS, but then you'll lose direct access from host...

Our host system is a CentOS with Kernel 2.6.16-026test015.2-smp

the guest systems are CentOS or Debian.

Hope you could help us, thanks TSchilli

Subject: Re: Fully access to NIC

Posted by [aistis](#) on Tue, 01 Aug 2006 11:17:35 GMT

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

try veth interface inside VPS:

http://wiki.openvz.org/Differences_between_venet_and_veth

<http://wiki.openvz.org/Veth>

Subject: Re: Fully access to NIC

Posted by [TSchilli](#) on Tue, 01 Aug 2006 14:00:35 GMT

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

Thanks for that pages.

I used the HOWTO to install such an VETH-device, but it doesn't work. The VPS have the MAC-Address given by the command, ping and everything works correct.

But the command arp still gives no reply, the arp-table seems to be empty or not readable.

Here my networkconfigs, hope somebody could find the mistake?

@ VPS

=====

```
-bash-3.00# ifconfig -a
```

```
eth0    Link encap:Ethernet  HWaddr 00:11:43:E3:F4:EF
        inet addr:192.168.0.101  Bcast:0.0.0.0  Mask:255.255.255.255
        inet6 addr: fe80::211:43ff:fee3:f4ef/64  Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
```

TX packets:6 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:384 (384.0 b)

lo Link encap:Local Loopback
inet addr:127.0.0.1 Mask:255.0.0.0
inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:127.0.0.1 P-t-P:127.0.0.1 Bcast:0.0.0.0 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
RX packets:4 errors:0 dropped:0 overruns:0 frame:0
TX packets:4 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:336 (336.0 b) TX bytes:336 (336.0 b)

venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:134.101.21.87 P-t-P:134.101.21.87 Bcast:134.101.21.87 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1

=====

@ VE0

=====

eth0 Link encap:Ethernet HWaddr 00:11:43:E3:F4:EE
inet addr:134.101.28.143 Bcast:134.101.31.255 Mask:255.255.240.0
inet6 addr: fe80::211:43ff:fee3:f4ee/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:12612560 errors:0 dropped:0 overruns:0 frame:0
TX packets:33275 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:2145404405 (1.9 GiB) TX bytes:4672347 (4.4 MiB)
Base address:0xecc0 Memory:dfae0000-dfb00000

eth1 Link encap:Ethernet HWaddr 00:11:43:E3:F4:EF
BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
Base address:0xdcc0 Memory:df8e0000-df900000

sit0 Link encap:IPv6-in-IPv4

NOARP MTU:1480 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
RX packets:9207 errors:0 dropped:0 overruns:0 frame:0
TX packets:3619 errors:0 dropped:16 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:1852791 (1.7 MiB) TX bytes:264012 (257.8 KiB)

veth102.0 Link encap:Ethernet HWaddr 00:11:43:E3:F4:EE
inet6 addr: fe80::211:43ff:fee3:f4ee/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:6 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:6 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:384 (384.0 b) TX bytes:0 (0.0 b)

Thanks TSchilli

Subject: Re: Fully access to NIC
Posted by [aistis](#) on Tue, 01 Aug 2006 20:11:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

TSchilli, put IP address from 192.168.0.0/24 subnet on veth102.0 on VE0 and try again.

Subject: Re: Fully access to NIC
Posted by [TSchilli](#) on Wed, 02 Aug 2006 06:11:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

With an address from 192.168.0./24 the guest is unreachable, I gave him an address from 134.101.16.0/20, our current netmask.
I'm able to ping the guest, or ping out of it. But arp still doesn't work.

Here the config again:

Host

=====

```
veth102.0 Link encap:Ethernet HWaddr 00:11:43:E3:F4:EE
  inet addr:134.101.21.87 Bcast:134.101.255.255 Mask:255.255.0.0
  inet6 addr: fe80::211:43ff:fee3:f4ee/64 Scope:Link
  UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
  RX packets:6 errors:0 dropped:0 overruns:0 frame:0
  TX packets:0 errors:0 dropped:6 overruns:0 carrier:0
  collisions:0 txqueuelen:0
  RX bytes:384 (384.0 b) TX bytes:0 (0.0 b)
```

=====

Guest

=====

```
eth0 Link encap:Ethernet HWaddr 00:11:43:E3:F4:EF
  inet addr:192.168.0.101 Bcast:0.0.0.0 Mask:255.255.255.255
  inet6 addr: fe80::211:43ff:fee3:f4ef/64 Scope:Link
  UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
  RX packets:0 errors:0 dropped:0 overruns:0 frame:0
  TX packets:6 errors:0 dropped:0 overruns:0 carrier:0
  collisions:0 txqueuelen:0
  RX bytes:0 (0.0 b) TX bytes:384 (384.0 b)
```

=====

Sorry, hopefully it helps....thanks

P.S.: Is it ok, that eth0 from guest has an 192.168.0./24 address while there is a different netmask at host?

Subject: Re: Fully access to NIC
Posted by [aistis](#) on Wed, 02 Aug 2006 10:17:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

it should be reachable with ip from that subnet if you have proper routing table. Look at wiki again, make sure you have everything in place.

Subject: Re: Fully access to NIC
Posted by [TSchilli](#) on Wed, 09 Aug 2006 08:10:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

I've checked all the settings, now I'm able to use the eth0-device in the VPS. Everything works fine. Tcpcdump and arp seems to work correctly.

The arp-table isn't empty anymore, but any any host has the mac address from our host system. Is that normal behavior, or something wrong with my configuration?

Thanks so far, TSchilli

Subject: Re: Fully access to NIC
Posted by [Vasily Tarasov](#) on Wed, 09 Aug 2006 08:31:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

Yeap, it's normal. All packets are routed from the host to the VE, so certainly arp can see only ARP-address of the host.

Subject: Re: Fully access to NIC
Posted by [TSchilli](#) on Wed, 09 Aug 2006 08:54:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

Is it possible to change this?
Some workarounds or something like that?

We want to use arpswatch in a VPS, seems a little bit stupid when we only get a single mac address for all the clients!

Thanks TSchilli

Subject: Re: Fully access to NIC
Posted by [Vasily Tarasov](#) on Wed, 09 Aug 2006 09:00:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

I see only one way: move network device to VE (--netdev_add).
But after that HOST system doesn't have this device.

Subject: Re: Fully access to NIC
Posted by [TSchilli](#) on Wed, 09 Aug 2006 09:05:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

OK, thanks for the great support TSchilli

P.S.: If somebody else have another idea please reply to this post or inform me via mail

Subject: Re: Fully access to NIC
Posted by [Andrey Mirkin](#) on Wed, 09 Aug 2006 09:29:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

There is another way to solve this problem.

You can create bridge in host system and add eth0 device and veth102.0 to it.

(See http://wiki.openvz.org/Virtual_Ethernet_device#Virtual_ether_net_devices_can_be_joined_in_one_bridge for details).

Do not forget to assign IP from eth0 host-system device to bridge device and reconfigure route table accordingly.

After that packets from outer network will be passed to VPS without routing in VE0, so you will see destination MAC-addresses in arp-table inside VPS.
