
Subject: VETH problems - USE DOM0 ip as a source [SOLVED]

Posted by [gucio1200](#) on Wed, 04 Jan 2012 09:21:42 GMT

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

Hi i have some problems to set cluster with ndb or galera - all the time is shows not container ip but containers host.

Lets say i have 2 server with proxmox using openvz

10.0.11.0/24 - its lan

10.0.1.0/16 - its private lan for each proxmox

it has

proxmox host#1	proxmox host#2
----------------	----------------

10.0.11.1-----	10.0.11.2
----------------	-----------

10.0.1.1	10.0.1.2
----------	----------

host#1	host#2
--------	--------

container(A)	container(B)
--------------	--------------

10.0.11.50	10.0.11.51
------------	------------

10.0.1.50	10.0.1.51
-----------	-----------

(gw 10.0.1.1)	(gw 10.0.1.2)
---------------	---------------

ndb cluster show somethink like that:

Failed to allocate nodeid, error: 'Error: Could not alloc node id at 10.0.11.50 port 1186: Connection with id 3 done from wrong host ip 10.0.11.1, expected 10.0.11.51(10.0.11.51).'

/etc/network/interaface on proxmox host#1 (on host#2 is simillar change only lan number from 1 to 2 of end ip):

where eth0 is gateway to net

```
auto eth0
```

```
iface eth0 inet static
```

```
address xxx.xxx.xxx.xxx
```

```
broadcast xxx.xxx.xxx.xxx
```

```
netmask xxx.xxx.xxx.xxx
```

```
gateway xxx.xxx.xxx.xxx
```

```
pointopoint xxx.xxx.xxx.xxx
```

```
auto vmbr0
```

```
allow-hotplug vmbr0
```

```
iface vmbr0 inet static
```

```
address 10.0.11.1
```

```
netmask 255.255.255.0
```

```
bridge_ports eth1
```

```
bridge_stp off
```

```
bridge_fd 0
```

```
auto vmbr1
allow-hotplug vmbr1
iface vmbr1 inet static
    address 10.0.1.1
    netmask 255.255.0.0
    bridge_ports tap0
    bridge_stp off
    bridge_fd 0
    pre-up /usr/sbin/openvpn --mktun --dev tap0
    pre-down /usr/sbin/openvpn --rmtun --dev tap0
```

AND MASQ:

```
10.0.11.0/24 on vmbr0 10.0.1.0/16
```

```
10.0.1.0/16 on vmbr1 10.0.11.0/24
```

sysctl.conf:

```
### Hetzner Online AG installimage
# sysctl config
net.ipv4.ip_forward=1
#net.ipv4.conf.all.rp_filter=1
#net.ipv4.icmp_echo_ignore_broadcasts=1
```

```
#net.ipv4.conf.default.forwarding=1
#net.ipv4.conf.default.proxy_arp = 1
#kernel.sysrq = 1
#net.ipv4.conf.default.send_redirects = 1
#net.ipv4.conf.all.send_redirects = 1
#net.ipv4.conf.eth0.proxy_arp=1
net.ipv4.ip_nonlocal_bind=1
```

```
net.ipv4.conf.default.forwarding = 1
net.ipv4.conf.all.forwarding = 1
net.ipv4.ip_forward=1
```

```
net.ipv4.conf.default.proxy_arp = 1
net.ipv4.conf.all.proxy_arp = 1
net.ipv4.conf.vmbr1.proxy_arp = 1
net.ipv4.conf.vmbr0.proxy_arp = 1
```

```
net.ipv4.conf.default.rp_filter = 0
net.ipv4.conf.all.rp_filter = 0
net.ipv4.conf.vmbr0.rp_filter = 0
net.ipv4.conf.vmbr1.rp_filter = 0
```

And VZCTL.conf

```
# On Hardware Node we generally need
# packet forwarding enabled and proxy arp disabled
```

```
net.ipv4.ip_forward = 1
net.ipv4.conf.default.forwarding=1
#net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.default.proxy_arp = 1

# Enables source route verification
net.ipv4.conf.all.rp_filter = 0

# Enables the magic-sysrq key
kernel.sysrq = 1

# TCP Explicit Congestion Notification
#net.ipv4.tcp_ecn = 0

# we do not want all our interfaces to send redirects
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0

net.ipv4.conf.default.rp_filter = 0
net.ipv4.conf.all.rp_filter = 0
net.ipv4.conf.eth0.rp_filter = 0
```

Has some one similar situation ?

To explain more My vms show me wrong ips source it looks like they connect from host machine. So lets say we have 10.0.11.50(vm1) on first host with openvz (10.0.11.1) and 10.0.11.51(vm2) on the second openvz (10.0.11.2)

openvz servers connected with crossover cable on eth1 bridged to vmbr1. I would like to connect from vm1 to vm2 and i can but in tcpdump i see they use 10.0.11.2 to connect. I will be very grateful for help. If you have any idea - just give me a tip or somethink becouse i dont have more ideas whats wrong ...

Should i use [hxxp://wiki.openvz.org/Virtual_Ethernet_device](http://wiki.openvz.org/Virtual_Ethernet_device) - "Using a directly routed IPv4 with virtual Ethernet device" ?

PROBLEM SOLVED - LOOK AT THE MASQ - well so easy but.. :)

Just remove masq from vmbr0 and vmbr1 then vm's see each other with right ip's. To build this kind of env. i had to create openvpn in ct not ct0 and then masq bridged eth1 (vmbr0) and eth0 (vmbr1). Cheers! Can close!