
Subject: 2 interfaces local network routing issue possible bug?

Posted by [QuantumNet](#) on Fri, 27 Jul 2007 05:09:31 GMT

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

Okay I have racked my head against the wall for hours... softlayer network engineers dont know enough about OpenVZ for this one so hopefully someone from here has gone through this already.

server has the typical softlayer interfaces eth0 and eth1

Openvz is installed and functioning (venet)

problem is the VPS' can communicate with the local interface even if they dont have a local IP assigned to them.

so lets say my eth0 IP was 10.21.1.3

and I assigned a VPS a public IP of 75.21.221.23 (routed to public vlan)

from this vps I can ping 10.21.1.3 even know it does not have a private IP assigned to it

Also if I assign another VPS a private IP address lets say 10.21.1.36

I can ping this one as well from the vps that does not have a private IP address.

so I did a traceroute from this cps and it shows it going through eth0 10.21.1.3 to get to 10.21.1.36

ip route list
(IPs altered)

```
10.21.1.128/26 dev eth0 proto kernel scope link src 10.21.1.3
169.254.0.0/16 dev eth0 scope link
10.0.0.0/8 via 10.21.1.2 dev eth0
```

I even tried deleting these routes and doing ip route flush cache

I am going nutz I dont want a VPS to have access to the hardware node's local IP address or any of the other VPS's local IP addresses unless I specifically assign a local IP to that VPS.

now mind you this only affects local IP's that are on the same machine if its on another server in the vlan then the VPS's dont have a route

God I hope someone knows the answer to this

Subject: Re: 2 interfaces local network routing issue possible bug?

Posted by [khorenko](#) on Fri, 27 Jul 2007 07:33:46 GMT

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

Hello.

1) well, might be it's not too bad that a VE with public IP can ping VEs/HN with local IPs? What problems can this bring?

2) the packets routing scheme is following: a packet from 75.21.221.23 goes to the VE0, there routing table is looked up for 10.21.1.3 destination, such a rule exists by default thus the packet is forwarded to venet.

Taking this into account if you still want to disable ability to ping VEs/HN with local IPs from a VE with global IP, you can create an iptables rule in VE0 which will drop packets with source - global IP of the mentioned VE and the local addresses 10.21.1.3/24 (or so) as a destination.

something like

```
# iptables -A FORWARD -s 75.21.221.23 -d 10.21.1.3/24 -j DROP
```

Hope this helps.

Subject: Re: 2 interfaces local network routing issue possible bug?

Posted by [QuantumNet](#) on Fri, 27 Jul 2007 13:36:55 GMT

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

It a bad thing because we use the backend network for out of band management, it is accessed via VPN then things like SSH, IPMI etc are available. The VPS' have no business being able to connect to it... it only provides a hole for them to be able to attack the private network.

So IPtables is the only way to curb this behavior huh? I figured there would have been a proper way to disable it in the kernel routing or some config file.

Subject: Re: 2 interfaces local network routing issue possible bug?

Posted by [khorenko](#) on Mon, 30 Jul 2007 11:12:48 GMT

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

QuantumNet wrote on Fri, 27 July 2007 17:36So IPtables is the only way to curb this behavior huh? I figured there would have been a proper way to disable it in the kernel routing or some config file.

Let's imagine you have 2 physical nodes A (public IP) and B (private IP) in the same segment. Can you ping the node B from A?

If you set the correct routing on the A (dev eth0) the answer will be YES if node B has a route back to A and no iptables rules on B prevent this.

In your case node A - the VE, node B - the Hardware Node. Node B already has a route to node A (or no packets at all will reach node A (VE)). Thus the only way to deny the ping from VE to HN is iptables.

Well, not exactly. You can completely change the network scheme: you can set up a bridge on the HN and use veth interface for VE instead of venet. Create and configure a bridge on the HN and add physical eth0 and corresponding veth interface to it. In this case you can remove route to the VE's IP on the HN and set up the exact default gateway (not just "dev" but also "via") on HN. Thus if that default gateway won't know about the VE's IP, the packets from HN won't reach VE.

But IMHO, iptables variant is much simpler if you definitely want to deny the connections.

Subject: Re: 2 interfaces local network routing issue possible bug?

Posted by [QuantumNet](#) on Mon, 30 Jul 2007 23:11:22 GMT

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

Quote:Let's imagine you have 2 physical nodes A (public IP) and B (private IP) in the same segment. Can you ping the node B from A?

If you set the correct routing on the A (dev eth0) the answer will be YES if node B has a route back to A and no iptables rules on B prevent this.

In your case node A - the VE, node B - the Hardware Node. Node B already has a route to node A (or no packets at all will reach node A (VE)). Thus the only way to deny the ping from VE to HN is iptables.

I dont agree with this one bit, it is not default behavior to make the public IP space communicate with the private IP space unless otherwise specified in NAT redirection or the system routing rules but it is not a default behavior it has to be later implimented.

Anyways I curbed this issue as stated with iptables:

```
iptables -A FORWARD -i venet0 -s 75.21.221.23 -d 10.21.1.3/24 -j DROP
iptables -A INPUT -i venet0 -s 75.21.221.23 -d 10.21.1.3/24 -j DROP
```

then in each VE that I wanted to be able to communicate with the backend network I assigned a local IP to it and set a corrusponding routing rule:

```
ip route add 10.0.0.0/8 dev venet0 src 10.21.1.32
```