
Subject: Not understanding networking (beware - tons of questions!) [SOLVED]

Posted by [jergendutch](#) on Wed, 30 Apr 2008 16:19:35 GMT

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

Hello,

I don't understand a lot about the way openvz handles networking.

I'm posting here because the forum seems really active, and hopefully someone can help

1. The wiki says that veth is dangerous because a container owner can forge mac addresses and ip addresses, but can't the host protect against this?

2. I am using Ubuntu inside a container, and I get a default gateway of 192.0.2.1. I have no idea where this comes from and I can't ping it from the container. Should openvz handle this automagically?

3. Should I use a default gateway of the host, or of my router inside the container? tcpdump shows the data leaving venet0, but I don't see how the data could ever get back.

4. How does the data get back?

5. There are two interfaces in the container, venet0 and venet0:0. The forum mentions this in a few places, but I can't find a post telling me why this exists. I've seen that there are various scripts for different Linux distributions but I can't see how openvz knows that a machine is e.g. Ubuntu. Is this the right direction to be going in? (At the moment I manually edit /etc/network/interfaces to remove the alias.)

Sorry for the mass of questions

Subject: Re: Not understanding networking (beware - tons of questions!)

Posted by [maratrus](#) on Mon, 05 May 2008 15:48:08 GMT

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

Hi,

sorry for delay

1.
Quote:
The wiki says that veth is dangerous because a container owner can forge mac addresses and ip addresses, but can't the host protect against this?

How can we change MAC address from inside the VPS? I supposed that it's impossible.

2.
Quote:

I am using Ubuntu inside a container, and I get a default gateway of 192.0.2.1. I have no idea where this comes from and I can't ping it from the container. Should openvz handle this automagically?

One of the /etc/vz/dists/scripts scripts adds this route. If you use simple venet configuration don't bother about this route. Our packets follow to the HN. But we cannot say that this route doesn't make any sense. We should have any default gateway inside VE otherwise neighbour table inside VE can be overflowed.

3.

Quote:

Should I use a default gateway of the host, or of my router inside the container? tcpdump shows the data leaving venet0, but I don't see how the data could ever get back

If you want to set your own configuration because of your network configurations you can change rc.local or other init scripts to put in order your network routes or something.

4.

Quote:

How does the data get back? Smile

Try to listen with "tcpdump" venet interface on HN.

5.

Quote:

There are two interfaces in the container, venet0 and venet0:0. The forum mentions this in a few places, but I can't find a post telling me why this exists. I've seen that there are various scripts for different Linux distributions but I can't see how openvz knows that a machine is e.g. Ubuntu. Is this the right direction to be going in? (At the moment I manually edit /etc/network/interfaces to remove the alias.)

Could you please describe the problem in more detail?

Subject: Re: Not understanding networking (beware - tons of questions!)

Posted by [jergendutch](#) on Tue, 06 May 2008 12:45:49 GMT

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

Thanks for your help

Can I check I've understood properly?

maratrus wrote on Mon, 05 May 2008 11:48

Hi,

sorry for delay

1.

Quote:

The wiki says that veth is dangerous because a container owner can forge mac addresses and ip addresses, but can't the host protect against this?

How can we change MAC address from inside the VPS? I supposed that it's impossible.

So have I understood correctly?

It is 100% impossible to change the MAC address inside the container for a veth device?

maratrus wrote on Mon, 05 May 2008 11:48

2.

Quote:

I am using Ubuntu inside a container, and I get a default gateway of 192.0.2.1. I have no idea where this comes from and I can't ping it from the container. Should openvz handle this automagically?

One of the `/etc/vz/dists/scripts` scripts adds this route. If you use simple venet configuration don't bother about this route. Our packets follow to the HN. But we cannot say that this route doesn't make nay sense. We should have any default gateway inside VE otherwise neighbour table inside VE can be overflowed.

So inside a VE with venet devices, the gateway is irrelevant - it can be anything?

maratrus wrote on Mon, 05 May 2008 11:48

3.

Quote:

Should I use a default gateway of the host, or of my router inside the container? tcpdump shows the data leaving venet0, but I don't see how the data could ever get back

If you want to set your own configuration because of your network configurations you can change `rc.local` or other init scripts to put in order your network routes or something.

4.

Quote:

How does the data get back? Smile

Try to listen with "tcpdump" venet interface on HN.

So the host listens for its own ip plus the ip of the VEs too?

maratrus wrote on Mon, 05 May 2008 11:48

5.

Quote:

There are two interfaces in the container, venet0 and venet0:0. The forum mentions this in a few places, but I can't find a post telling me why this exists. I've seen that there are various scripts for different Linux distributions but I can't see how openvz knows that a machine is e.g. Ubuntu. Is this the right direction to be going in? (At the moment I manually edit /etc/network/interfaces to remove the alias.)

Could you please describe the problem in more detail?

Yes, ifconfig shows a device venet0 and a second device venet0.0 in the VE. Why is this? I am using the official Debian VE image with a CentOS host.

Subject: Re: Not understanding networking (beware - tons of questions!)

Posted by [maratrus](#) on Wed, 07 May 2008 11:50:44 GMT

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

Hello,

1.

Quote:

It is 100% impossible to change the MAC address inside the container for a veth device?

I have to admit that you've confused me with this question. So I've started to investigate this issue and found the following facts:

a) This feature was added here

"<http://git.openvz.org/?p=linux-2.6.18-openvz;a=commit;h=9b04401d0719c5a7f45d0816d3633a68>

ef79f10c"

b) But if you want to change MAC address from inside the VPS you have to set appropriate permission for this VE from inside the HN

c) The common methods of vzctl utility cannot do this.

http://bugzilla.openvz.org/show_bug.cgi?id=687

But you can write simple program for this purpose and give your VE such permission. If you want I can give you the simple example of such program.

2.

Quote:

So inside a VE with venet devices, the gateway is irrelevant - it can be anything?

yes

3.

Quote:

So the host listens for its own ip plus the ip of the VEs too?

a) vzctl put appropriate record to the route table when VE starts(something like VE_IP dev venet0 scope link)

b) When you try to reach your VE arp-request comes to HN.

The node will arp reply for VE_IP if and only if "ip r g VE_IP from HN_IP" (HN_IP form [incoming dev])

will return a route other than one to [incoming dev]

c) vzctl also put appropriate record to arp table to allow HN only reply for VE_IP addresses.

d) HN reply for the apr_request

c) Then HN receives other packets

HN works like a common router I'd say.

Subject: Re: Not understanding networking (beware - tons of questions!)

Posted by [jergendutch](#) on Wed, 07 May 2008 12:38:12 GMT

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

Thanks very much - very, very useful
