Subject: *UPDATED* networking Problem venet / veth
Posted by dasicebaer on Mon, 12 Feb 2007 17:30:27 GMT
View Forum Message <> Reply to Message

Hi, there!

I'm currently setting up a server with ovz028-test015.1-patched kernel and vz*-tools version 3.0.11 with Debian Etch as base for VE0 and VEs. There'll be two VEs, one for the internal network named lanserv (supplying samba, dhcp, etc) and one with the dmz-applications named dmzserv (www, mail, dns).

From what I learned from the wiki, I need veth for programs like samba or dhcp due to the need for incoming broadcasts, which has the disadvantage of being less secure than venet. So I would like to have venet for dmzserv and veth for lanserv.

However, lanserv automatically gets a venet0 at startup in its /etc/network/interfaces. There is no interfaces.template on my system (mentioned in the file), nor a {VEID}.start-file to edit (solved a problem in a previous post) nor do I want to completely disable networking in the /etc/init.d/vz-script (also from a previous post). Any hints where to differentiate the network virtualizations between the VEs?

I would also like to automatically add the appropriate routes on start of lanserv and do not know where (beside building a custom script for that). Anyone already solved that one?

At last I'd like to ask if anyone else has problems of very long lockups (up to three minutes) when starting or shutting down VEs with more than one ipaddress, which I encounter at the moment (might be related to the above problem of having both venet and veth-devices in lanserv). cpu and memusage are completely okay during that time, no unusuably high load from the harddrives either...

Thanks for openvz and for answering in advance!

Subject: Re: networking Problem venet / veth
Posted by Vasily Tarasov on Tue, 13 Feb 2007 15:42:51 GMT
View Forum Message <> Reply to Message

Hello,

First of all I want to say, that veth is less secure only in untrusted environments. The thing is that many people use OpenVZ for hosting: they sell VEs to various people. In such case this is rather insecure to give veth device to VE, because it will be able to create any ethernet packets, can be sniffered and etc.

Secondly, why venet0 in lanserv bothers you? Let it be there, but not configured. It will not prevent veth somehow!

As concerns appropriate routes for veth. There is a file /usr/sbin/vznetcfg for it. It is a bash script,

which is invoked by vzctl on VE start. Look at its source and you'll understand what to do in order to add appropriate routes.

HTH,
Vasily.

Subject: Re: networking Problem venet / veth
Posted by dasicebaer on Thu, 15 Feb 2007 10:46:29 GMT
View Forum Message <> Reply to Message

Vasily Tarasov wrote on Tue, 13 February 2007 10:42Hello,

First of all I want to say, that veth is less secure only in untrusted environments. The thing is that many people use OpenVZ for hosting: they sell VEs to various people. In such case this is rather insecure to give veth device to VE, because it will be able to create any ethernet packets, can be sniffered and etc.

Well, this is why I don't want to use veth for the dmzserv. From a security aspect, it's not a good idea to have dmzserv and lanserv running on the same machine at all, since the dmz's only purpose is to confine an attacker in it's own lansegment, where he can't do any harm to the internal network. I hope to achieve this confinement through some firewallrules, but if the attacker could simply change his ip to move the dmzserv from the dmz to the inner lan, that would render the dmz pretty useless.

Quote:Secondly, why venet0 in lanserv bothers you? Let it be there, but not configured. It will not prevent veth somehow!

Okay. First of all, I recognized the debianpackage of vzutils from etch did not contain /usr/sbin/vznetcfg. After reinstalling from source, I did not experience any more lags where I had to wait up to five minutes after a "vzctl enter VEID".

Secondly, I don't know how to specify that the ipaddress will be bound to the vethdevice and not to the venet. My configuration is like this:

NAME="lanserv"
VETH="veth2.0,00:00:00:00:00:02,eth0,00:00:00:00:00:22"
IP_ADDRESS="172.17.1.242"

After starting the VE, vzlist will show me the correct ipaddress (but not the hostname btw), but neiter VE nor VE0 have the correct routes added:

VE0:/etc/vz/conf# ip route ls
172.17.1.242 dev venet0 scope link [...]

lanserv:/# ip route ls
lanserv:/# ifconfig -a
eth0   Link encap:Ethernet HWaddr 00:00:00:00:00:22

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:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

lo Link encap:Local Loopback
LOOPBACK 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
BROADCAST POINTOPOINT NOARP 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:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

Basically, I have to go through the complete procedure in
http://wiki.openvz.org/Virtual_Ethernet_device#Simple_config
uration_with_virtual_ethernet_device again in VE and VE0 to be able to use networking of the VE.
Is there any option I forgot to give in the VEID.conf-file or anything?

Quote:As concerns appropriate routes for veth. There is a file /usr/sbin/vznetcfg for it. It is a bash script, which is invoked by vzctl on VE start. Look at its source and you'll understand what to do in order to add appropriate routes.

My vznetcfg doesn't do anything beside an "ifconfig ${dev} up". So I added veth2.0 to VE0's /etc/network/interfaces, but it wouldn't help - still the same as quoted above:

VE0:/etc/network/# cat interfaces
[...]iface veth2.0 inet static
up sysctl -w net.ipv4.conf.veth2.0.forwarding=1
up sysctl -w net.ipv4.conf.vetn2.0.proxy_arp=1
up route add 172.17.1.242 dev veth2.0

Any more ideas? If not, I would start writing some bashscripts to run via vznetcfg for adding the appropriate rules to VE0 / some init.d-script for the VEs.

Thanks for helping so far!

Subject: Re: networking Problem venet / veth
Posted by Vasily Tarasov on Thu, 15 Feb 2007 11:39:38 GMT
View Forum Message <> Reply to Message

Quote:
Secondly, I don't know how to specify that the ipaddress will be bound to the vethdevice and not to the venet. My configuration is like this:

NAME="lanserv"
VETH="veth2.0,00:00:00:00:00:02,eth0,00:00:00:00:00:22"
IP_ADDRESS="172.17.1.242"

After starting the VE, vzlist will show me the correct ipaddress (but not the hostname btw), but neiter VE nor VE0 have the correct routes added:

VE0:/etc/vz/conf# ip route ls
172.17.1.242 dev venet0 scope link [...]

IP_ADDRESS parameter is only for venet interface! So don't use it for VE, that has only veth interface.

In order to produce veth interface configuration automatically the following algorithm should be used, I guess:

vzctl set 101 --netif_add eth0,00:12:34:56:78:9A,veth101.0,00:12:34:56:78:9B --save
Produce eth0 configuration inside VE by means, that are template-specific. I mean, you should add appropriate IP/route information in certain configuration files. For example, in gentoo there is a file: /etc/conf.d/net. If you add there
config_eth0=("192.168.0.2/24")
eth0 interface in VE will be automatically configured. In a similar manner the routing configuration in VE should be made.
In the last vzctl version vznecfg script is able to call external script $EXTERNAL_SCRIPT, that should be mentioned in /etc/vz/vznet.conf. In this script you should tune appropriate routing on VE0.

Actually a Wiki-page should be written on this subject. May be you? We'll ve appreciate to you

Thanks,
HTH,
Vasily

Subject: Re: networking Problem venet / veth
Posted by dasicebaer on Thu, 15 Feb 2007 11:57:30 GMT
View Forum Message <> Reply to Message

Vasily Tarasov wrote on Thu, 15 February 2007 12:39

Produce eth0 configuration inside VE by means, that are template-specific. I mean, you should add appropriate IP/route information in certain configuration files. For example, in gentoo there is

a file: /etc/conf.d/net. If you add there
config_eth0=("192.168.0.2/24")
eth0 interface in VE will be automatically configured. In a similar manner the routing configuration in VE should be made.

okay, in debian there's /etc/network/interfaces, where all interfaces are defined. This one gets overwritten by /etc/vz/conf/dists/scripts/debian-add_ip.sh. Since the interfaces.template-file is included before the dynamincally created interfacesentries, I can't remove interfaces. I notice, that the script doesn't care about veth at all, so it should be rewritten (yup, I'll do that now).

Could you just clarify when the script is called in the init-process?

Quote:
In the last vzctl version vznecfg script is able to call external script $EXTERNAL_SCRIPT, that should be mentioned in /etc/vz/vznet.conf. In this script you should tune appropriate routing on VE0.

Actually a Wiki-page should be written on this subject. May be you? We'll ve appreciate to you

I'll happily do so, as soon as I got everything working.

Subject: Re: networking Problem venet / veth
Posted by Vasily Tarasov on Thu, 15 Feb 2007 12:23:29 GMT
View Forum Message <> Reply to Message

Quote:
okay, in debian there's /etc/network/interfaces, where all interfaces are defined. This one gets overwritten by /etc/vz/conf/dists/scripts/debian-add_ip.sh. Since the interfaces.template-file is included before the dynamincally created interfacesentries, I can't remove interfaces. I notice, that the script doesn't care about veth at all, so it should be rewritten (yup, I'll do that now).

Hmm... But, as far as I understand /etc/vz/conf/dists/scripts/debian-add_ip.sh script should work only if IP_ADDRESS is in <veid>.conf file! Does it run even if IP_ADDRESS is missing? In your case, when VE has only veth active , IP_ADDRESS should be missing.

Quote:
Could you just clarify when the script is called in the init-process?

Which one? /etc/vz/conf/dists/scripts/debian-add_ip.sh? It runs before init process...

Quote:
I'll happily do so, as soon as I got everything working.

Thanks!

Subject: Re: networking Problem venet / veth
Posted by dasicebaer on Thu, 15 Feb 2007 13:02:11 GMT
View Forum Message <> Reply to Message

so we would have to define the ipaddress for veth-enabled VEs somewhere else? hrm... I guess it would be better if the init-script could differentiate between venet and veth-ips. Shouldn't be too hard to define in the configfile.

Is there any possible benefit or scenario, in which a VE should have both interfacetypes, veth and venet? If not, the VEID.conf:IP_ADDRESS-entry could be the same for both and the init-script would call either add_ip_venet.sh or add_ip_veth.sh, based on a switch if NETIF is set or not.

Else there would have to be two entries for the ip...

I'll see how much time I can spare on this one, work is calling for now. Thanks for the help so far!

Subject: Re: networking Problem venet / veth
Posted by Vasily Tarasov on Thu, 15 Feb 2007 13:10:00 GMT
View Forum Message <> Reply to Message

I completely agree with you, that not all is ok here. And in case, when both interface types (veth and venet) should be utilized in VE, the situation is more harder.

You can actually fill the bug and some decision will be made.

Good luck in your work!

Subject: Updated WIKI
Posted by dasicebaer on Fri, 16 Feb 2007 13:58:21 GMT
View Forum Message <> Reply to Message

Alright, the problem seems only solveable by changing the vzctl-sources which is too much for me at the moment. So I wrote a workaround in the wiki, which should leave VENETs intact while allowing VETHs. Not perfect, but works.

Please comment if you have problems in your configuration.