Hello,

Since my local tests with OpenVZ were working perfectly, I've installed yesterday a few boxes in production running OpenVZ in order to deploy my services and to replace my current Xen architecture.

So far so good.... up to the network part... and I'm seriously about to hang myself, I've been stuck for like 7 hours straight.... so any kind of help would be GREATLY appreciated

Here's the deal :

I want to reproduce a similar network topology I had with Xen which was to put all the HN (dom0) inside a private subnet 10.1.0.0/16, outside of Internet, with a gateway 10.1.1.1 providing internet access to them, to save a few IPs and to isolate the HNs from public access. And to give to the VEs (domU), IPs from a public routable subnet provided by my ISP, and the VEs to directly connect to the gateway of the ISP (their router) as their default route.

With Xen it wasn't a problem, I just had to create an interface for the domU, to set it up like a normal box directly connected to the Internet with its public IP and the ISP gateway, and that's all....

From what I've read through the documentation and the wiki, it seems it is not possible in that way with OpenVZ... at least not with the venet device....

Am I right ?

I was quite disappointed, until I found the veth device in the wiki.... which seems to be the real thing for me and to act the way that I want

[I don't care about the possible security issue, since I am not a reseller dealing with evil customers, and I'm running in a fully trusted environment]

But still, it's been a few hours that I've tried to make the thing work and I can't :/

I've followed the examples here http://wiki.openvz.org/Virtual\_Ethernet\_device

So far, I have the veth101.0 device on the HN for my VE with ID 101, I also have the eth0 on the VE.

I've put all the sysctl calls explained on the wiki, the route add on both the HN and the VE, etc... but still... all I can do is :

ping the VE from the HN with it's Internet IP. ping the HN and its subnet

and nothing more (no one can ping the VE from the outside except the HN, and I don't have Internet from the VE)

Is what I want to do impossible with the current way OpenVZ is built ? Has someone succeeded ? If yes, is the wiki complete ? And someone could share its configuration if it has a similar setup running ?

Thanks a lot in advance !

Ugo

Subject: Re: VEs with different subnets Posted by dev on Mon, 02 Jul 2007 07:29:55 GMT View Forum Message <> Reply to Message

have you read the following article aboute multiple devices/GW and routed venet: http://wiki.openvz.org/Source\_based\_routing ?

The typical configuration of secured OVZ looks the following way:

HN VEs

eth0 (assigned local IP address)

eth1 (assigned global IP address) VEs are routed and set as default default GW through eth1, since it's default GW.

the only difference is that eth1 still has global IP assigned to HN. But this is a plus on the other hand, since HN often needs an access to internet for example for upgrades.

In this case you simple need to put "eth1" to /etc/vz/conf variable: 17 # The name of the device whose ip address will be used as source ip for VE. 18 # By default automatically assigned. 19 #VE\_ROUTE\_SRC\_DEV="eth1"

Subject: Re: VEs with different subnets Posted by dev on Mon, 02 Jul 2007 07:38:02 GMT View Forum Message <> Reply to Message Quote:

Is what I want to do impossible with the current way OpenVZ is built ?

I forgot to add: sure it IS \*possible\* if you don't want to assign IP to eth1, then bridges are your choice.

if give an access we can help you in configuring.

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 02 Jul 2007 08:31:58 GMT View Forum Message <> Reply to Message

dev wrote on Mon, 02 July 2007 03:29have you read the following article aboute multiple devices/GW and routed venet: http://wiki.openvz.org/Source\_based\_routing ?

The typical configuration of secured OVZ looks the following way:

HN VEs

eth0 (assigned local IP address)

eth1 (assigned global IP address) VEs are routed and set as default default GW through eth1, since it's default GW.

the only difference is that eth1 still has global IP assigned to HN. But this is a plus on the other hand, since HN often needs an access to internet for example for upgrades.

In this case you simple need to put "eth1" to /etc/vz/conf variable: 17 # The name of the device whose ip address will be used as source ip for VE. 18 # By default automatically assigned. 19 #VE\_ROUTE\_SRC\_DEV="eth1"

Hello,

Thanks for you answer.

Yes, I did read the wiki "Source based routing", but it implies to have a new device set up for each new route that I want to give access to, to my VEs, right ? So it means to assign a global IP to each HN ?

As I said, I would like to avoid "wasting" global IPs on the HNs, because I only have 64 IP and about 20 HNs... so it would mean only 44 IPs left for the VEs, which is not a lot in my case. That's why I've built the explained configuration with Xen... well and of course the HN (dom0) have Internet access from their private subnet, because 10.1.1.1 acts as a NAT Gateway to them, so I can do upgrades, etc... with the HNs.

Ugo

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 02 Jul 2007 08:36:57 GMT View Forum Message <> Reply to Message

dev wrote on Mon, 02 July 2007 03:38Quote: Is what I want to do impossible with the current way OpenVZ is built ?

I forgot to add: sure it IS \*possible\* if you don't want to assign IP to eth1, then bridges are your choice.

if give an access we can help you in configuring.

Yes, that's what I've figured.... (bridges = veth, right ?).... the veth looks like more what I want to do....

The wiki configuration on veth is not enough ? Do you know any sources where I could find more configuration examples on it ?

But yes of course if I don't get it, I will provide you an access... but don't want to abuse your time.

Ugo

Subject: Re: VEs with different subnets Posted by dev on Mon, 02 Jul 2007 08:52:27 GMT View Forum Message <> Reply to Message

1. veth != bridge. veth are ethernet like devices which act as a tunnel. one veth device is in VE0 and another one in VE.

After this tunnel is configured one need to setup how veth end in VE0 traffic is routed or bridged.

2. bridges are used to stack together a number of eth devices (including veth). So one need to add real eth adapter to bridge and all the desired veth devices in VE0. brctl tool is used for this.

3. I still believe your case can be easily done with venet and routing. Will check in a minute and return back to you.

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 02 Jul 2007 09:13:11 GMT View Forum Message <> Reply to Message

dev wrote on Mon, 02 July 2007 04:521. veth != bridge. veth are ethernet like devices which act as a tunnel. one veth device is in VE0 and another one in VE.

After this tunnel is configured one need to setup how veth end in VE0 traffic is routed or bridged.

2. bridges are used to stack together a number of eth devices (including veth). So one need to add real eth adapter to bridge and all the desired veth devices in VE0. brctl tool is used for this.

3. I still believe your case can be easily done with venet and routing. Will check in a minute and return back to you.

1/ Ohh ok 2/ Hmm, ok.. 3/ You would save my day

Thanks.

Ugo

Subject: Re: VEs with different subnets Posted by khorenko on Mon, 02 Jul 2007 11:32:16 GMT View Forum Message <> Reply to Message

Hello, Ugo.

just a small question - did you try to use venet and just configure node as simple as you did on your Xen node?

i mean just set a private IP on a Hardware Node (eth0) and set a public IP for a VE? (vzctl set \$VEID --ipadd \$PUBLIC\_IP --save)

i just checked such a simple configuration and it seems to me - it works fine. Could you please check it? If it won't work, could you please post

- ip a l
- ip r l
- vzctl exec \$VEID ip a l
- vzctl exec \$VEID ip r l
- the command that didn't work?

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 02 Jul 2007 12:01:18 GMT View Forum Message <> Reply to Message

finist wrote on Mon, 02 July 2007 07:32Hello, Ugo.

just a small question - did you try to use venet and just configure node as simple as you did on your Xen node?

i mean just set a private IP on a Hardware Node (eth0) and set a public IP for a VE? (vzctl set \$VEID --ipadd \$PUBLIC\_IP --save)

i just checked such a simple configuration and it seems to me - it works fine.

Could you please check it? If it won't work, could you please post

- ip a l
- ip r l
- vzctl exec \$VEID ip a l
- vzctl exec \$VEID ip r l
- the command that didn't work?

Yes of course, that's how I've tested OpenVZ at home and everything worked fine (but I was on the same private subnet).

And of course, I've tried the same thing at the datacenter with the public IP directly.... but it failed.

Just give me 10 minutes, the time for me to delete everything and to create a fresh new VE and everything.... and I will post everything that you want.

I'll create a Fedora VE this time, because the Debian VE (from the templates provided by openvz.org) doesn't have the ip-tools suite.... and since I don't have Internet... I can't apt-get install it :/

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 02 Jul 2007 12:19:46 GMT View Forum Message <> Reply to Message

titan:/var/lib/vz/template/cache# vzctl create 101 --ostemplate fedora-core-6-i686-default --config vps.basic

Creating VE private area (fedora-core-6-i686-default) Performing postcreate actions VE private area was created titan:/var/lib/vz/template/cache# vzctl set 101 --onboot yes --save Saved parameters for VE 101 titan:/var/lib/vz/template/cache# vzctl set 101 --hostname test --save Saved parameters for VE 101 titan:/var/lib/vz/template/cache# vzctl set 101 --ipadd 87.98.196.135 --save Saved parameters for VE 101 titan:/var/lib/vz/template/cache# vzctl set 101 --ipadd 87.98.196.135 --save

that's how I've created the VZ (just in case).

87.98.196.135, being the public IP address I want to try to assign to the VE. Here are the parameters that would do the trick if the VE had a "classical" ethernet device : address 87.98.196.135 netmask 255.255.255.192 network 87.98.196.128 broadcast 87.98.196.191 gateway 87.98.196.129

Here's the output of the commands you've requested. (titan is the HN, test the VE)

titan:~# ip a l

2: bond0: <broadcast,multicast,master> mtu 1500 qdisc noop</broadcast,multicast,master>
link/ether 00:00:00:00:00 brd ff:ff:ff:ff:ff
4: lo: <loopback,up,10000> mtu 16436 qdisc noqueue</loopback,up,10000>
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127.0.0.1/8 scope host lo
inet6 ::1/128 scope host
valid_lft forever preferred_lft forever
6: eth0: <broadcast,multicast,up,10000> mtu 1500 qdisc pfifo_fast qlen 1000</broadcast,multicast,up,10000>
link/ether 00:50:70:26:53:9d brd ff:ff:ff:ff:ff
inet 10.1.1.15/16 brd 10.1.255.255 scope global eth0
inet6 fe80::250:70ff:fe26:539d/64 scope link
valid_lft forever preferred_lft forever
8: dummy0: <broadcast,noarp> mtu 1500 qdisc noop</broadcast,noarp>
link/ether 12:62:b8:e1:94:93 brd ff:ff:ff:ff:ff
10: teql0: <noarp> mtu 1500 qdisc noop qlen 100</noarp>
link/void
12: tunl0: <noarp> mtu 1480 qdisc noop</noarp>
link/ipip 0.0.0.0 brd 0.0.0.0
14: gre0: <noarp> mtu 1476 qdisc noop</noarp>
link/gre 0.0.0.0 brd 0.0.0.0
16: sit0: <noarp> mtu 1480 qdisc noop</noarp>
link/sit 0.0.0.0 brd 0.0.0.0
18: ip6tnl0: <noarp> mtu 1460 qdisc noop</noarp>
link/tunnel6 00:00:00:00:00:00:00:00:00:00:00:00:00:

titan:~# ip r l 87.98.196.135 dev venet0 scope link 10.1.0.0/16 dev eth0 proto kernel scope link src 10.1.1.15 default via 10.1.1.1 dev eth0 titan:~#

(10.1.1.1 being my gateway providing NAT Internet access for the HNs, and not the ISP one)

titan:~# vzctl exec 101 ip a l

1: lo: <LOOPBACK,UP,LOWER\_UP> mtu 16436 qdisc noqueue link/loopback 00:00:00:00:00 brd 00:00:00:00:00:00 inet 127.0.0.1/8 scope host lo inet6 ::1/128 scope host valid\_lft forever preferred\_lft forever

3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER\_UP> mtu 1500 qdisc noqueue link/void inet 127.0.0.1/32 scope host venet0 inet 87.98.196.135/32 brd 87.98.196.135 scope global venet0:0 titan:~#

titan:~# vzctl exec 101 ip r l 192.0.2.0/24 dev venet0 scope host 169.254.0.0/16 dev venet0 scope link default via 192.0.2.1 dev venet0 titan:~#

(I really don't get the 192.0.2.0/24 that I've just seen from the route....)

I can ping from the HN and the VE can ping the HN

titan:~# ping 87.98.196.135 PING 87.98.196.135 (87.98.196.135) 56(84) bytes of data. 64 bytes from 87.98.196.135: icmp\_seq=1 ttl=64 time=0.063 ms 64 bytes from 87.98.196.135: icmp\_seq=2 ttl=64 time=0.017 ms

[root@test /]# ping 10.1.1.15 PING 10.1.1.15 (10.1.1.15) 56(84) bytes of data. 64 bytes from 10.1.1.15: icmp\_seq=1 ttl=64 time=0.061 ms 64 bytes from 10.1.1.15: icmp\_seq=2 ttl=64 time=0.017 ms But nothing else, but it's quite logical since I can't set the gateway.... and I don't know how to do it with the venet device....

Just in case, here's the version of the vzctl program

titan:~# vzctl vzctl version 3.0.16-5dso1 Copyright (C) 2000-2007 SWsoft. This program may be distributed under the terms of the GNU GPL License.

Thanks a lot.

Ugo

Subject: Re: VEs with different subnets Posted by khorenko on Mon, 02 Jul 2007 16:08:58 GMT View Forum Message <> Reply to Message

Thank you for the logs. yes, my suggestion implied that 10.1.1.1 will route the packets from VEs also as it does for Hardware Nodes.

Well, if you want to set up another gateway for VEs then just try source base routing as dev already suggested.

# /sbin/ip rule add from 87.98.196.135 table 10# /sbin/ip route add default dev eth0 via 87.98.196.129 table 10

This should do the trick i suppose. (you won't be forced to do this manually each time, when you get the commands consequence required to get it work, you can create a script with these commands).

If it won't work again, please, let me know, we'll think about the reasons. Or just provide an access to the node, it will take much less time to undestand and eliminate the problem.

ugo123 wrote on Mon, 02 July 2007 16:19titan:~# vzctl exec 101 ip r l 192.0.2.0/24 dev venet0 scope host 169.254.0.0/16 dev venet0 scope link default via 192.0.2.1 dev venet0 titan:~#

(I really don't get the 192.0.2.0/24 that I've just seen from the route....)

This IP 192.0.2.1 is fake one, it's not used, venet is just a point-to-point conenction, so all the packets got to the venet0 inside a VE appeared on the venet0 interface on a Hardware Node and are routed further according to the rules on Hardware Node.

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 02 Jul 2007 17:17:06 GMT View Forum Message <> Reply to Message

finist wrote on Mon, 02 July 2007 12:08Thank you for the logs. yes, my suggestion implied that 10.1.1.1 will route the packets from VEs also as it does for Hardware Nodes.

Well, if you want to set up another gateway for VEs then just try source base routing as dev already suggested.

# /sbin/ip rule add from 87.98.196.135 table 10# /sbin/ip route add default dev eth0 via 87.98.196.129 table 10

This should do the trick i suppose. (you won't be forced to do this manually each time, when you get the commands consequence required to get it work, you can create a script with these commands).

If it won't work again, please, let me know, we'll think about the reasons. Or just provide an access to the node, it will take much less time to undestand and eliminate the problem.

ugo123 wrote on Mon, 02 July 2007 16:19titan:~# vzctl exec 101 ip r l 192.0.2.0/24 dev venet0 scope host 169.254.0.0/16 dev venet0 scope link default via 192.0.2.1 dev venet0 titan:~#

(I really don't get the 192.0.2.0/24 that I've just seen from the route....)

This IP 192.0.2.1 is fake one, it's not used, venet is just a point-to-point conenction, so all the packets got to the venet0 inside a VE appeared on the venet0 interface on a Hardware Node and are routed further according to the rules on Hardware Node.

Ok, thank you both again for you support (at least with Xen I sure had the network, but not the support and the people behind) hahaha

Yes I don't want to use 10.1.1.1 as the gateway for the VE....

Because it would be a Single Point Of Failure for the whole network.... whereas the gateway of my ISP is an huge backbone and is fully redondant in many ways...

But it's perfectly acceptable for providing Internet access to the HN (just for the upgrades)

I'm 99% sure that I've tried what you're suggesting, and that it blocked the network of the HN

where I've tried that. (as I told dev, I've already read and tried the Source Based Routing)

But I will try again and tell you if it works or not with your EXACT same lines.... but I actually can't test it right now because I've locked the network of one of the HN by wrongfully trying to add the veth device and the eth device to a same bridge (it locked the network too).

So now I need to go to the datacenter in order to reboot all of my hard failures I will tell you ASAP (in about 2h) if it works or not with the Source Based Routing.

Thanks again.

Ugo

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 02 Jul 2007 20:13:38 GMT View Forum Message <> Reply to Message

Ok, now I remember why source routing wasn't working :

mercury:~# /sbin/ip rule add from 87.98.196.137 table 10 mercury:~# /sbin/ip route add default dev eth0 via 87.98.196.129 table 10 RTNETLINK answers: Network is unreachable mercury:~#

mercury:~# ping 87.98.196.129

PING 87.98.196.129 (87.98.196.129) 56(84) bytes of data. 64 bytes from 87.98.196.129: icmp\_seq=1 ttl=254 time=0.845 ms 64 bytes from 87.98.196.129: icmp\_seq=2 ttl=254 time=0.654 ms 64 bytes from 87.98.196.129: icmp\_seq=3 ttl=254 time=0.724 ms

--- 87.98.196.129 ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 1998ms rtt min/avg/max/mdev = 0.654/0.741/0.845/0.078 ms mercury:~#

mercury:~# vzctl exec 101 ip a l

 lo: <LOOPBACK,UP,LOWER\_UP> mtu 16436 qdisc noqueue link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00 inet 127.0.0.1/8 scope host lo inet6 ::1/128 scope host valid\_lft forever preferred\_lft forever
venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER\_UP> mtu 1500 qdisc noqueue link/void inet 127.0.0.1/32 scope host venet0 inet 87.98.196.137/32 brd 87.98.196.137 scope global venet0:0

mercury:~#

(yes I've done this on another HN and another IP for the VE, but this is EXACTLY the same network, configuration, etc...)

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 02 Jul 2007 21:19:39 GMT View Forum Message <> Reply to Message

And I have to remember that I only have one interface and my goal is actually not to waste a public IP on the HN since they don't need one at all, they just need an internet access for the upgrade, and that's where my NAT gateway 10.1.1.1 takes place to save some IPs.

After reading the route manual, that's probably why I've got the "RTNETLINK answers: Network is unreachable" message because the program doesn't see any DIRECT interface that knows how to reach the ISP's real gateway.

Actually, after thinking about it.... I don't see how this is do-able at all with "classical routing", because even if we succeed, the ISP will surely filters the "martians" source IP (10.1.x.x) from its routing table, so the packets will be null-routed anyway... and the thing will be pointless in that case

So that's why I think that I really need the veth to get the thing working...because from what I see the veth looks like a lot like the things I was able to do with Xen networking part.... but still I can't get it work (at least by by doing the things explained on the wiki)

Thanks,

Ugo

Subject: Re: VEs with different subnets Posted by khorenko on Tue, 03 Jul 2007 07:20:58 GMT View Forum Message <> Reply to Message

ugo123 wrote on Tue, 03 July 2007 00:13Ok, now I remember why source routing wasn't working :

mercury:~# /sbin/ip rule add from 87.98.196.137 table 10 mercury:~# /sbin/ip route add default dev eth0 via 87.98.196.129 table 10 RTNETLINK answers: Network is unreachable mercury:~#

This is simply means that Hardware Node does not know how to get to 87.98.196.129. Could you

please try again with:

# ip route add 87.98.196.128/26 dev eth0# ip rule add from 87.98.196.137 table 10# ip route add default dev eth0 via 87.98.196.129 table 10

This should eliminate the "Network is unreachable" error.

Subject: Re: VEs with different subnets Posted by ugo123 on Tue, 03 Jul 2007 09:25:30 GMT View Forum Message <> Reply to Message

finist wrote on Tue, 03 July 2007 03:20ugo123 wrote on Tue, 03 July 2007 00:13Ok, now I remember why source routing wasn't working :

mercury:~# /sbin/ip rule add from 87.98.196.137 table 10 mercury:~# /sbin/ip route add default dev eth0 via 87.98.196.129 table 10 RTNETLINK answers: Network is unreachable mercury:~#

This is simply means that Hardware Node does not know how to get to 87.98.196.129. Could you please try again with:

# ip route add 87.98.196.128/26 dev eth0# ip rule add from 87.98.196.137 table 10# ip route add default dev eth0 via 87.98.196.129 table 10

This should eliminate the "Network is unreachable" error.

Thanks.

Yes, that's what I've figured in my other message.... (and I forgot to mention that I've tried the ip route add 87.98.196.128/26 after some Google-ing).

But still it hasn't worked .... :/

And as I posted... I think that the problem is that the ISP's gateway is null-routing the packets coming from private subnets (martians IPs) as this is a common security rule.

Ugo

Subject: Re: VEs with different subnets Posted by ugo123 on Tue, 03 Jul 2007 09:26:36 GMT

## View Forum Message <> Reply to Message

mercury:~# ip r l 87.98.196.137 dev venet0 scope link src 10.1.1.14 87.98.196.128/26 dev eth0 scope link 10.1.0.0/16 dev eth0 proto kernel scope link src 10.1.1.14 default via 10.1.1.1 dev eth0 mercury:~# ip a l 2: bond0: <BROADCAST,MULTICAST,MASTER> mtu 1500 qdisc noop link/ether 00:00:00:00:00 brd ff:ff:ff:ff:ff:ff 4: Io: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00 inet 127.0.0.1/8 scope host lo inet6 ::1/128 scope host valid Ift forever preferred Ift forever 6: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo\_fast qlen 1000 link/ether 00:50:70:26:57:37 brd ff:ff:ff:ff:ff:ff inet 10.1.1.14/16 brd 10.1.255.255 scope global eth0 inet6 fe80::250:70ff:fe26:5737/64 scope link valid\_lft forever preferred\_lft forever 8: dummy0: <BROADCAST,NOARP> mtu 1500 gdisc noop link/ether e6:1f:aa:72:10:de brd ff:ff:ff:ff:ff:ff:ff 10: tegl0: <NOARP> mtu 1500 gdisc noop glen 100 link/void 12: tunl0: <NOARP> mtu 1480 qdisc noop link/ipip 0.0.0.0 brd 0.0.0.0 14: gre0: <NOARP> mtu 1476 qdisc noop link/gre 0.0.0.0 brd 0.0.0.0 16: sit0: <NOARP> mtu 1480 qdisc noop link/sit 0.0.0.0 brd 0.0.0.0 18: ip6tnl0: <NOARP> mtu 1460 gdisc noop 1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 gdisc nogueue link/void mercury:~#

## Subject: Re: VEs with different subnets Posted by khorenko on Tue, 03 Jul 2007 10:28:44 GMT View Forum Message <> Reply to Message

Is it possible to get an access to the node? Please, it will be much more simple... Just for security sake: I confirm that "finist" (Konstantin Khorenko) is OpenVZ developer. His email is khorenko@sw.ru :@)

Subject: Re: VEs with different subnets Posted by ugo123 on Tue, 03 Jul 2007 10:41:12 GMT View Forum Message <> Reply to Message

Ok, no problem....

I'm going to prepare everything for your access and to mail you with the according information.

Thanks,

Ugo

Subject: Re: VEs with different subnets Posted by emkravts on Tue, 03 Jul 2007 11:45:06 GMT View Forum Message <> Reply to Message

Hello, Ugo.

I've set up configuration you described and it works. As a result wiki page appeared devoted to subj. Please check

http://wiki.openvz.org/Using\_veth\_and\_brctl\_for\_protecting\_H N\_and\_saving\_IP-adresses

Regards Evgeny.

Subject: Re: VEs with different subnets Posted by ugo123 on Tue, 03 Jul 2007 12:04:17 GMT View Forum Message <> Reply to Message

dev wrote on Tue, 03 July 2007 06:31 Just for security sake: I confirm that "finist" (Konstantin Khorenko) is OpenVZ developer. His email is khorenko@sw.ru :@)

No problems

I've just sent an email to Konstantin to its private email.

Thanks again.

Ugo

Subject: Re: VEs with different subnets Posted by ugo123 on Tue, 03 Jul 2007 12:07:26 GMT View Forum Message <> Reply to Message

emkravts wrote on Tue, 03 July 2007 07:45Hello, Ugo.

I've set up configuration you described and it works. As a result wiki page appeared devoted to subj. Please check

http://wiki.openvz.org/Using\_veth\_and\_brctl\_for\_protecting\_H N\_and\_saving\_IP-adresses

Regards Evgeny.

Oh, thanks a lot !

I'm checking this right now.... so you confirm that I need to use veth in order to do that ? (as I was expecting it -- it's not a problem for me)

Ugo

Subject: Re: VEs with different subnets Posted by dev on Tue, 03 Jul 2007 12:08:10 GMT View Forum Message <> Reply to Message

let's try via venet with Kostya first

Subject: Re: VEs with different subnets Posted by ugo123 on Tue, 03 Jul 2007 12:39:28 GMT View Forum Message <> Reply to Message

dev wrote on Tue, 03 July 2007 08:08let's try via venet with Kostya first

Hehe, internal competition between OpenVZ guys ? Well, you have full access now to try

I'm doing meanwhile (on another box, don't worry) the configuration Evgeny explained, and which is closer to the Xen network configuration I had.

emkravts wrote on Tue, 03 July 2007 07:45Hello, Ugo.

I've set up configuration you described and it works. As a result wiki page appeared devoted to subj. Please check

http://wiki.openvz.org/Using\_veth\_and\_brctl\_for\_protecting\_H N\_and\_saving\_IP-adresses

Regards Evgeny.

Hmmm, it really seems like the exact way I was thinking about it (except that during my tests I was trying everything on a single interface eth0... that was my mistake... I'm a bridge newbie...), but there's something I don't understand, it is at the first step....

eth1 Link encap:Ethernet HWaddr XX:XX:XX:XX:XX:35 inet addr:192.168.0.32 Bcast:192.168.3.255 Mask:255.255.252.0 inet6 addr: fe80::213:d4ff:fe90:4d50/64 Scope:Link UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1 RX packets:603734 errors:0 dropped:0 overruns:0 frame:0 TX packets:36627 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000 RX bytes:0 (0.0 b) TX bytes:0 (0.0 b) Interrupt:21

Why does eth1 in your example has an IP ? shouldn't it be IP-less ?

I guess it's something you've forgot when you have copy/paste, but if it's not the case, I don't get it :/

And do you know if this configuration could work with a single interface eth0 (I only have one physical ethernet interface for these servers), but with an alias interface instead of a second real one (something like eth0:0 instead of eth1)?

Thanks again

Ugo

Subject: Re: VEs with different subnets Posted by emkravts on Tue, 03 Jul 2007 13:09:23 GMT View Forum Message <> Reply to Message

We are removing IP from eth1 in item 4). But if you have OpenVZ kernel + tools installed on the box - you do not need IP for eth1 at the beginning. As for eth0 - it is only to make HN accessible

from LAN. If you have only one eth device - it should be treated as eth without IP for connecting to ISP (instead of eth1 on the picture). We should just erase eth0+LAN from picture and read eth1 as eth0

Subject: Re: VEs with different subnets Posted by ugo123 on Tue, 03 Jul 2007 14:48:36 GMT View Forum Message <> Reply to Message

emkravts wrote on Tue, 03 July 2007 09:09We are removing IP from eth1 in item 4). But if you have OpenVZ kernel + tools installed on the box - you do not need IP for eth1 at the beginning. As for eth0 - it is only to make HN accessible from LAN. If you have only one eth device - it should be treated as eth without IP for connecting to ISP (instead of eth1 on the picture). We should just erase eth0+LAN from picture and read eth1 as eth0

Hmm, but I still need to connect to the LAN, despite the fact that I have only one interface (I still need to connect to the HN from time to time, to upgrade, etc...)

And I can't bring an alias without an IP (I've never tried this before)

pandora:/var/lib/vz/template/cache# /sbin/ifconfig eth0:0 0 SIOCSIFFLAGS: Cannot assign requested address pandora:/var/lib/vz/template/cache#

Shit... that was so close :/ Damn cheap servers with only one eth interface.

Subject: Re: VEs with different subnets Posted by n00b\_admin on Mon, 09 Jul 2007 12:35:03 GMT View Forum Message <> Reply to Message

I know i'm not much of an expert but i'm learning too from your experience...

Why do you need public addresses on the VE's ?

You assign public addresses to the HN's and private to the VE's.. what you need to do that cannot be done because of NAT-ing the VE's ?

You setup the VE's with venet0 and that's all you need to do.

Subject: Re: VEs with different subnets Posted by ugo123 on Mon, 09 Jul 2007 12:51:19 GMT View Forum Message <> Reply to Message n00b\_admin wrote on Mon, 09 July 2007 08:35I know i'm not much of an expert but i'm learning too from your experience...

Why do you need public addresses on the VE's ?

You assign public addresses to the HN's and private to the VE's.. what you need to do that cannot be done because of NAT-ing the VE's ?

You setup the VE's with venet0 and that's all you need to do.

Hello,

I need public addresses on the VEs because those VEs are the one who are providing Internet services, and I want one public IP per VE.

I don't want to assign a single public IP to the HN because it's a waste of IP (and being in Europe, we are tight on public IP with the RIPE), it's also useless in my case and it exposes the HN to the Internet (even with a firewall, I don't like the idea).

The HN should be for me the most secured box (because if compromised, everything is going down), and a private subnet is ideal to answer both of my problems : no public IP wasted, impossible to reach from the Internet.

I don't want to do NAT either, because NAT is less than ideal both in terms of performance and configuration, it would be of course okay if I had a single HN and a single IP. like set the port 25 to my mail\_ve, 80 to my web\_ve, etc..etc..

But it's not my case and I want a full network capacity on each VE... and to configure each VE the most easy way, like a real box.... and to don't mind any above configuration.... like NATing and so on... so if a HN dies, I can migrate the VE to the HN, launch it again, and it directly works... no configuration involved.

Finally I could have tweaked my main gateway 10.1.1.1 for my internal network to provide a kind of mixed routing to support my case, but it would have meant that the WHOLE network would have relied and transited on a single machine aka a SPoF (Single Point of Failure), whereas my ISP is providing me a nice IP gateway, with full redondancy, heavy reliability, etc..etc....

So I guess for most simple cases or when you can't trust your VE, venet is definitively the way to go....

But when you need to create more complex infrastructure, it has some limitations, it has nothing to do with the way OpenVZ is built, it's just the limitation of Layer 3 IP routing.... you sometimes need to a Level 2 (with MAC addresses and so on) to do more tricky things.

Hope it answers your questions.

Ugo

Your VE's provide internet access to other private networks or what ?

I manage a server where i use the setup i'm talking about, because i live in Europe too

I don't think that exposing a HN that is running only the ssh daemon to internet is a great security risk. Maybe i'm wrong but i'm using such setup for four months now and i didn't have any incidents regarding the HN. Besides the normal dictionary attacks on the ssh daemon which i block using public key auth ( a normal thing to do in these days ).

On this setup i provide internet services for several domains and as needed will increase the domains hosted on the HN.

Right now i provide DNS, Mail, Database, web and ftp access to 10 domains without any problem whatsoever.

For two production boxes that i run using openvz (the other one hosts 40+ domains) i'm pretty happy on how the software performs and how secure it is.

I'm running openvz on the 40+ domains box from last years spring and i didn't had any break-in's on the HN itself. That tells me that the kernel is secure enough to not allow a compromised VE access to the HN. I had compromised VE's but mostly kids vandalizing sites because they were poorly written.

On that production box i'm running the VE's with public ip's but the incidents were only related to php applications that were not sanitized properly and after the developers fixed the problems saw by me i stopped having problems with that too.

It's your choice on how paranoid you want to be with security and if you have sensitive data your handling it's normal to be that way but if you're only a web hoster i think it's to much of a hassle to do this setup your trying to do.

Subject: Re: VEs with different subnets Posted by khorenko on Sun, 15 Jul 2007 16:31:56 GMT View Forum Message <> Reply to Message

BTW, this is solved and a wiki page was created describing this configuration:

http://wiki.openvz.org/Using\_private\_IPs\_for\_Hardware\_Nodes