
Subject: OpenVZ problem non stop
Posted by [maser](#) on Sat, 26 Apr 2008 12:55:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

hi there, I have been using openvz for about a week now, everything was running just fine when i had it installed, now suddenly all virtual servers are unaccessable and you cant ping/ssh to them, while the host is okay. I can still be able to ssh from the box to them, but not from outside, I tried to bind the ips to the host then unbind them and stop and start the vds's and they worked well for an hour then again suddenly all goes offline, you cant access from or to them. It's really strange having them off and on I tried to check my sysctl.conf and nothing wrong with it compared to your guide, they only worked when I used the centos script in network to bind all the virtual server ips to the host then unbind them back, it works for an hour then again goes all offline, I would appreciate it so much if you can help me as I got 22 virtual servers running for my customers and it's so bad to get that up and down and unstability, thank you.

Subject: Re: OpenVZ problem non stop
Posted by [maratrus](#) on Mon, 28 Apr 2008 09:39:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

could you please describe your problem in more detail.

1. What's your network configuration? Do you use venet or veth?

- ip a | (from HN)
- ip ro | (from HN)
- ip rule | (from HN)
- arp -n (from HN when you lose an access to CTs)
- sysctl -a | grep forward (from HN when it stops working)

2. Anything strange in logs or in dmesg?

3. You can use "tcpdump" utility to explore packets behavior.

When it stops working could you please try to ping one of your CTs from external machine. In that moment try to use "tcpdump" utility on the external machine and on the HN. What does it output?

Thank You!

Subject: Re: OpenVZ problem non stop
Posted by [jyrppa](#) on Mon, 28 Apr 2008 16:12:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

I have similar problem. My VPS obtain IP address via DHCP. I can ping and ssh my VPS some time and then suddenly I can't. TCPDUMP shows incoming messages to but no replies back. There is no resource shortages or any other error logs. Sometime when I can't connect VPS I still can reach main hardware node.

When the problem is on and I go to VPS locally and give dhclient command it sometimes helps..

Subject: Re: OpenVZ problem non stop

Posted by [jyrppa](#) on Tue, 29 Apr 2008 18:51:09 GMT

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

- ip a l (from HN):

2: lo: <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_lft forever preferred_lft forever

4: sit0: <NOARP> mtu 1480 qdisc noop

link/sit 0.0.0.0 brd 0.0.0.0

6: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000

link/ether 00:0d:9d:96:dd:54 brd ff:ff:ff:ff:ff:ff

inet 212.54.16.183/20 brd 212.54.31.255 scope global eth0

inet6 fe80::20d:9dff:fe96:dd54/64 scope link

valid_lft forever preferred_lft forever

8: br0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue

link/ether 00:0d:9d:96:dd:54 brd ff:ff:ff:ff:ff:ff

inet 212.54.16.183/20 brd 212.54.31.255 scope global br0

inet6 fe80::20d:9dff:fe96:dd54/64 scope link

valid_lft forever preferred_lft forever

1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue

link/void

3: veth101.0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue

link/ether 00:0d:9d:96:dd:54 brd ff:ff:ff:ff:ff:ff

inet6 fe80::20d:9dff:fe96:dd54/64 scope link

valid_lft forever preferred_lft forever

- ip ro l (from HN):

212.54.16.0/20 dev br0 proto kernel scope link src 212.54.16.183

212.54.16.0/20 dev eth0 proto kernel scope link src 212.54.16.183

default via 212.54.16.1 dev eth0

default via 212.54.16.1 dev br0

- ip rule l (from HN):

0: from all lookup 255

32766: from all lookup main

32767: from all lookup default

- arp -n (from HN when you lose an access to CTs)

Address	HWtype	HWaddress	Flags	Mask	Iface
212.54.16.1	ether	00:05:31:F8:B6:39	C		br0

- sysctl -a | grep forward (from HN when it stops working):

```
error: "Operation not permitted" reading key "net.ipv6.route.flush"
error: "Operation not permitted" reading key "net.ipv4.route.flush"
net.ipv6.conf.veth101/0.forwarding = 0
net.ipv6.conf.eth0.forwarding = 0
net.ipv6.conf.br0.forwarding = 0
net.ipv6.conf.default.forwarding = 0
net.ipv6.conf.all.forwarding = 0
net.ipv6.conf.lo.forwarding = 0
net.ipv4.conf.veth101/0.mc_forwarding = 0
net.ipv4.conf.veth101/0.forwarding = 1
net.ipv4.conf.venet0.mc_forwarding = 0
net.ipv4.conf.venet0.forwarding = 1
net.ipv4.conf.br0.mc_forwarding = 0
net.ipv4.conf.br0.forwarding = 1
net.ipv4.conf.eth0.mc_forwarding = 0
net.ipv4.conf.eth0.forwarding = 1
net.ipv4.conf.lo.mc_forwarding = 0
net.ipv4.conf.lo.forwarding = 0
net.ipv4.conf.default.mc_forwarding = 0
net.ipv4.conf.default.forwarding = 1
net.ipv4.conf.all.mc_forwarding = 0
net.ipv4.conf.all.forwarding = 0
net.ipv4.ip_forward = 0
```

2. no

3.
from VPS:
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on venet0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
21:46:18.001065 IP e212-54-19-183.elisa-laajakaista.fi > e212-54-12-63.elisa-laajakaista.fi: ICMP
echo reply, id 41459, seq 213, length 64
21:46:19.012210 IP e212-54-19-183.elisa-laajakaista.fi > e212-54-12-63.elisa-laajakaista.fi: ICMP
echo reply, id 41459, seq 214, length 64
21:46:19.999945 IP e212-54-19-183.elisa-laajakaista.fi > e212-54-12-63.elisa-laajakaista.fi: ICMP
echo reply, id 41459, seq 215, length 64

from HN:
only ICMP echo requests can be seen but no replies.

So the echo replies are lost between VPS and HN.

Subject: Re: OpenVZ problem non stop
Posted by [maratrus](#) on Wed, 30 Apr 2008 09:38:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

Thank You for the information!

Could you possibly add some info please:

1. "ip a l" (from inside the VE)
2. "ip ro l" (from inside the VE)
3. "ip ru le" (from inside the VE)
4. When you try to ping your VE from the external node please:
"tcpdump -n -i br0" (on NH)
At the same time
"tcpdump -n -i eth0" (inside the VE) (I guess you have got eth0 interface inside VE)

P.S. You should set an IP address only to br0 interface not to the eth0 interface.

Subject: Re: OpenVZ problem non stop
Posted by [jyrppa](#) on Thu, 01 May 2008 06:06:30 GMT
[View Forum Message](#) <> [Reply to Message](#)

Here is the additional info:

```
# ip a l
1: lo: <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_lft forever preferred_lft forever
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue
    link/void
    inet 127.0.0.1/32 scope host venet0
5: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:0b:cd:fc:19:ff brd ff:ff:ff:ff:ff:ff
    inet 212.54.19.183/20 brd 212.54.31.255 scope global eth0
```

```
# ip ro l
192.0.2.1 dev venet0 scope link
212.54.16.0/20 dev eth0 proto kernel scope link src 212.54.19.183
default via 212.54.16.1 dev eth0
default via 192.0.2.1 dev venet0
```

```
# ip rule l
0:    from all lookup 255
32766: from all lookup main
32767: from all lookup default
```

(from ve)

```
# tcpdump -n -i eth0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
08:57:37.210148 IP 212.54.8.1.67 > 255.255.255.255.68: BOOTP/DHCP, Reply, length 317
08:57:37.212672 IP 212.54.8.1.67 > 255.255.255.255.68: BOOTP/DHCP, Reply, length 317
08:57:37.300555 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 27130, seq 43, length 64
08:57:38.301405 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 27130, seq 44, length 64
08:57:39.301362 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 27130, seq 45, length 64
08:57:40.219645 IP 212.54.8.1.67 > 255.255.255.255.68: BOOTP/DHCP, Reply, length 316
08:57:40.302579 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 27130, seq 46, length 64
```

(from ve)

```
# tcpdump -i venet0
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on venet0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
09:01:42.397499 IP e212-54-19-183.elisa-laajakaista.fi > e212-54-12-63.elisa-laajakaista.fi: ICMP
echo reply, id 41969, seq 41, length 64
```

(from HN)

```
#tcpdump -n -i br0
09:03:00.403286 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 41969, seq 119, length 64
09:03:01.406447 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 41969, seq 120, length 64
09:03:02.402553 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 41969, seq 121, length 64
09:03:03.400700 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 41969, seq 122, length 64
09:03:04.400841 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 41969, seq 123, length 64
09:03:05.406213 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 41969, seq 124, length 64
09:03:06.402854 IP 212.54.12.63 > 212.54.19.183: ICMP echo request, id 41969, seq 125, length 64
```

Subject: Re: OpenVZ problem non stop
Posted by [jyrppa](#) on Thu, 01 May 2008 06:10:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

>P.S. You should set an IP address only to br0 interface not to the eth0 interface.

Yes, now the configuration looks like:

```
2: lo: <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_lft forever preferred_lft forever
4: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
6: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:0d:9d:96:dd:54 brd ff:ff:ff:ff:ff:ff
    inet6 fe80::20d:9dff:fe96:dd54/64 scope link
        valid_lft forever preferred_lft forever
8: br0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:0d:9d:96:dd:54 brd ff:ff:ff:ff:ff:ff
    inet 212.54.16.183/20 brd 212.54.31.255 scope global br0
    inet6 fe80::20d:9dff:fe96:dd54/64 scope link
        valid_lft forever preferred_lft forever
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue
    link/void
3: veth101.0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:0d:9d:96:dd:54 brd ff:ff:ff:ff:ff:ff
    inet6 fe80::20d:9dff:fe96:dd54/64 scope link
        valid_lft forever preferred_lft forever
```

same problem still.

Subject: Re: OpenVZ problem non stop
Posted by [maratrus](#) on Thu, 01 May 2008 09:07:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

try to remove

default via 192.0.2.1 dev venet0

inside VE.

Subject: Re: OpenVZ problem non stop
Posted by [jyrppa](#) on Mon, 05 May 2008 09:15:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks a lot maratrus! Your solution worked.

Subject: Re: OpenVZ problem non stop
Posted by [l4ndy74](#) on Tue, 06 May 2008 10:57:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

why br0 and no eth0?
maratrus any ideas? http://forum.openvz.org/index.php?t=tree&goto=29845&#msg_29845

Subject: Re: OpenVZ problem non stop
Posted by [maratrus](#) on Wed, 07 May 2008 09:29:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

could you please describe your problem in more detail.

1. What's your network configuration? Do you use venet or veth?

- ip a | (from HN and from inside the VPS)
- ip ro | (from HN and from inside the VPS)
- ip rule | (from HN and from inside the VPS)
- arp -n (from HN when you lose an access to CTs)
- sysctl -a | grep forward (from HN when it stops working)

2. Anything strange in logs or in dmesg?

3. You can use "tcpdump" utility to explore packets behavior.

When it stops working could you please try to ping one of your CTs from external machine. In that moment try to use "tcpdump" utility on the external machine and on the HN. What does it output?

Quote:

why br0 and no eth0?

Sorry, I don't understand you?

br0 - it's a bridge

Thank You!
