

---

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.

---