
Subject: set alternative router for VE through dhcp
Posted by [eisenhorn](#) on Fri, 13 Mar 2009 21:34:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

I have ubuntu (8.10 with kernel from 8.04 supporting openv) box (HN) with 2 NIC: eth0 - ISP1, eth1 - LAN (192.168.7.1).

It acts as simple NAT (iptables -A POSTROUTING -s 192.168.7.0/24 -j MASQUERADE), so ip_forward is enabled.

On eth1 I created bridge (br0) to let VE see the network.

Also this box acts as DHCP-server.

On the other hand I have adsl-modem configured as router (192.168.7.3, ISP2). I'm trying to pass all traffic from some machines through ISP2. It is easily done by specifying "option routers 192.168.7.3;" in dhcpd.conf per machine. And it works fine with physical boxes, but not with VE (192.168.7.6).

In VE I'm trying to add default route, but it has no effect - no packets are going through router:

```
root@gretchin:/# ping ya.ru
PING ya.ru (213.180.204.8) 56(84) bytes of data.
```

```
--- ya.ru ping statistics ---
3 packets transmitted, 0 received, 100% packet loss, time 2007ms
```

```
root@gretchin:/# traceroute ya.ru
traceroute to ya.ru (213.180.204.8), 30 hops max, 40 byte packets
 1  192.168.7.3 (192.168.7.3)  2.059 ms  2.693 ms  2.997 ms
 2  * * *
```

I'm sure that it isn't router's issue, cause doing such on HN helps:

```
ip rule add from 192.168.7.6 table 6
ip route add default dev br0 via 192.168.7.3 table 6
```

But why it isn't running without adding routing rules on HN?

Subject: Re: set alternative router for VE through dhcp
Posted by [maratrus](#) on Mon, 16 Mar 2009 13:31:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

May be I've missed something but here is my opinion.

In spite of the fact that VE looks like a common physical server it has its own peculiarities.
A record that looks like

default via 192.168.7.3 dev eth0

on standalone physical server makes network packet pass through 192.168.7.3 and in our case the next hop would be adsl-modem.

But in case of VE network packet always pass through HN.

So, could you please confirm that VEs veth interface is united with eth1 into br0 on the HN?

Subject: Re: set alternative router for VE through dhcp
Posted by [eisenhorn](#) on Wed, 18 Mar 2009 09:35:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

I have these settings in /etc/network/interfaces:

```
auto lo
iface lo inet loopback

#ISP1
auto eth0
iface eth0 inet dhcp

#LAN (physical eth1)
auto br0
iface br0 inet static
    address 192.168.7.1
    netmask 255.255.255.0
    bridge_ports eth1
```

And veth's are adding dynamically (as described in wiki - through vznet.conf).
So "brctl show br0" displays

```
bridge name bridge id STP enabled interfaces
br0 8000.00024494a7ac no eth1
    veth101.0
```

Subject: Re: set alternative router for VE through dhcp
Posted by [maratrus](#) on Wed, 18 Mar 2009 09:45:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

It doesn't clear from this output if br0 contains anything except eth1.

Could you please also show tcpdump output while pinging from inside the VE (tcpdump output

from HN and from VE).

Subject: Re: set alternative router for VE through dhcp
Posted by [eisenhorn](#) on Wed, 18 Mar 2009 10:15:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

br0 has eth1 and veth101.0 - it's clear from "brctl show br0" I provided in previous post.

This is what is done when default GW points to 192.168.7.1 (HN, router with ISP1 on eth0):

```
listening on br0, link-type EN10MB (Ethernet), capture size 96 bytes
12:07:24.991494 IP 192.168.7.6 > ya.ru: ICMP echo request, id 9848, seq 259, length 64
12:07:25.026797 IP ya.ru > 192.168.7.6: ICMP echo reply, id 9848, seq 259, length 64
12:07:25.027327 IP 192.168.7.6.54562 > 192.168.7.1.domain: 27719+ PTR?
8.204.180.213.in-addr.arpa. (44)
12:07:25.027735 IP 192.168.7.1.domain > 192.168.7.6.54562: 27719 1/0/0 (63)
12:07:25.991493 IP 192.168.7.6 > ya.ru: ICMP echo request, id 9848, seq 260, length 64
12:07:26.029356 IP ya.ru > 192.168.7.6: ICMP echo reply, id 9848, seq 260, length 64
12:07:26.029879 IP 192.168.7.6.40815 > 192.168.7.1.domain: 42993+ PTR?
8.204.180.213.in-addr.arpa. (44)
12:07:26.030277 IP 192.168.7.1.domain > 192.168.7.6.40815: 42993 1/0/0 (63)
12:07:26.990485 IP 192.168.7.6 > ya.ru: ICMP echo request, id 9848, seq 261, length 64
12:07:27.026324 IP ya.ru > 192.168.7.6: ICMP echo reply, id 9848, seq 261, length 64
```

Doing "route add default gw 192.168.7.3" inside VE to bind it on ISP2 gives next results:

```
12:08:08.895129 IP 192.168.7.6.47667 > 192.168.7.1.domain: 38422+ A? ya.ru. (23)
12:08:08.899668 IP 192.168.7.1.domain > 192.168.7.6.47667: 38422 1/0/0 A ya.ru (39)
12:08:08.901466 arp who-has 192.168.7.3 tell 192.168.7.6
12:08:08.902118 arp reply 192.168.7.3 is-at 00:1c:f0:28:7b:3c (oui Unknown)
12:08:08.902194 IP 192.168.7.6 > ya.ru: ICMP echo request, id 13688, seq 1, length 64
12:08:09.901461 IP 192.168.7.6 > ya.ru: ICMP echo request, id 13688, seq 2, length 64
12:08:10.901514 IP 192.168.7.6 > ya.ru: ICMP echo request, id 13688, seq 3, length 64
12:08:11.901496 IP 192.168.7.6 > ya.ru: ICMP echo request, id 13688, seq 4, length 64
12:08:12.901496 IP 192.168.7.6 > ya.ru: ICMP echo request, id 13688, seq 5, length 64
12:08:13.901499 IP 192.168.7.6 > ya.ru: ICMP echo request, id 13688, seq 6, length 64
12:08:14.161418 arp who-has 192.168.7.6 tell 192.168.7.1
12:08:14.161465 arp reply 192.168.7.6 is-at 00:18:51:bb:fa:69 (oui Unknown)
12:08:14.901493 IP 192.168.7.6 > ya.ru: ICMP echo request, id 13688, seq 7, length 64
12:08:15.901503 IP 192.168.7.6 > ya.ru: ICMP echo request, id 13688, seq 8, length 64
```

Subject: Re: set alternative router for VE through dhcp
Posted by [maratrus](#) on Wed, 18 Mar 2009 11:00:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ok,

00:1c:f0:28:7b:3c belongs to adls-modem, doesn't it?

And does 00:18:51:bb:fa:69 belong to veth101.0?

Could you run tcpdump with "-e" parameter on all interfaces, i.e. on et1, br0, veth101.0 (on the HN) and on eth0 (inside VE, just in case)?

Subject: Re: set alternative router for VE through dhcp
Posted by [eisenhorn](#) on Wed, 18 Mar 2009 11:50:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

HN:

eth0 (ISP1): 00:02:44:8f:89:8f

eth1 (as well as br0): 00:02:44:94:a7:ac

veth101.0: 00:18:51:80:8a:d9

VE:

eth0: 00:18:51:bb:fa:69

modem:

eth0: 00:1C:F0:28:7B:3C

Will provide tcpdump bit later.
