
Subject: Not seeing multicast on veth

Posted by [zajdee](#) on Thu, 15 Oct 2009 09:27:40 GMT

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

Hello,

I have a problem with multicast not coming 'out' of the virtual machine.

I am trying to set up a test environment of a three virtual machines on a single OpenVZ master host. I have quite some experience with OpenVZ, however haven't found a single working solution for the problem I am seeing now.

The environment is as follows:

- virtual machines with Debian Lenny Linux, two veth interfaces
- master host with Debian Lenny Linux, two VLANs, each VLAN bridged to one of the veth interfaces of a virtual machine

The virtual machines run the Linux VRRP userspace daemon, ucarp, which uses multicast to communicate with other VMs.

And finally, the problem is, that I cannot see the multicast traffic coming out of the VM, if I run `tcpdump -i veth101.84` on master (101 is the VM's ID)

I can, however, see the traffic, if I run `tcpdump` inside the VM (`tcpdump -i eth0.84`).

The Linux kernel is 2.6.18-openvz el5.028stab064.7.szn5, based on RHEL's kernel el5.028stab064.7 with some specific addons for our environment.

I know that there has been an issue in the el5.028stab033 RHEL kernel, however that issue (according to various forums) has been resolved in stab034.

Does anyone has any idea, what could be the problem of not seeing the multicasts on veth?

Thanks in advance.

And now some details from the master:

`$ route -n`

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
10.2.12.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0.12
192.168.0.0	10.2.12.1	255.255.0.0	UG	0	0	0	eth0.12
172.16.0.0	10.2.12.1	255.240.0.0	UG	0	0	0	eth0.12
10.0.0.0	10.2.12.1	255.0.0.0	UG	0	0	0	eth0.12
0.0.0.0	10.2.12.1	0.0.0.0	UG	0	0	0	eth0.12

`$ brctl show`

bridge name	bridge id	STP enabled	interfaces
br0	8000.00238b89da94	no	eth0
br0.131	8000.0018511abf2b	no	eth0.131

		veth101.131
		veth102.131
		veth103.131
br0.84	8000.0018511abf2b	no eth0.84
		veth101.84
		veth102.84
		veth103.84

```
$ tcpdump -i veth101.84 '(vrrp)' -n
tcpdump: WARNING: veth101.84: no IPv4 address assigned
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on veth101.84, link-type EN10MB (Ethernet), capture size 96 bytes
11:23:05.927984 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120,
authtype none, intvl 1s, length 20
11:23:06.875967 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120,
authtype none, intvl 1s, length 20
11:23:06.876338 IP 10.2.84.62 > 224.0.0.18: VRRPv2, Advertisement, vrid 62, prio 0, authtype
none, intvl 5s, length 36
11:23:06.876360 IP 10.2.84.63 > 224.0.0.18: VRRPv2, Advertisement, vrid 63, prio 0, authtype
none, intvl 5s, length 36
11:23:06.876382 IP 10.2.84.61 > 224.0.0.18: VRRPv2, Advertisement, vrid 61, prio 0, authtype
none, intvl 5s, length 36
11:23:07.799989 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120,
authtype none, intvl 1s, length 20
11:23:08.748021 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120,
authtype none, intvl 1s, length 20
11:23:09.687968 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120,
authtype none, intvl 1s, length 20
11:23:10.679980 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120,
authtype none, intvl 1s, length 20
```

and also from the VM, the tcpdump was captured at the same time:

```
$ tcpdump -i eth0.84 '(vrrp)' -n
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0.84, link-type EN10MB (Ethernet), capture size 96 bytes
11:23:06.875967 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120,
authtype none, intvl 1s, length 20
11:23:06.876338 IP 10.2.84.62 > 224.0.0.18: VRRPv2, Advertisement, vrid 62, prio 0, authtype
none, intvl 5s, length 36
11:23:06.876360 IP 10.2.84.63 > 224.0.0.18: VRRPv2, Advertisement, vrid 63, prio 0, authtype
none, intvl 5s, length 36
11:23:06.876382 IP 10.2.84.61 > 224.0.0.18: VRRPv2, Advertisement, vrid 61, prio 0, authtype
none, intvl 5s, length 36
11:23:07.799989 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120,
authtype none, intvl 1s, length 20
11:23:07.800050 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 136, prio 0, authtype
```

none, intvl 1s, length 36
11:23:07.800053 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 138, prio 196, authtype none, intvl 1s, length 36
11:23:07.800057 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 137, prio 128, authtype none, intvl 1s, length 36
11:23:08.748021 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120, authtype none, intvl 1s, length 20
11:23:09.687968 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120, authtype none, intvl 1s, length 20
11:23:09.688028 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 138, prio 196, authtype none, intvl 1s, length 36
11:23:09.688030 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 137, prio 128, authtype none, intvl 1s, length 36
11:23:09.688032 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 136, prio 0, authtype none, intvl 1s, length 36
11:23:10.679980 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120, authtype none, intvl 1s, length 20
11:23:11.551986 IP 10.2.84.240 > 224.0.0.18: VRRPv2, Advertisement, vrid 84, prio 120, authtype none, intvl 1s, length 20
11:23:11.552036 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 136, prio 0, authtype none, intvl 1s, length 36
11:23:11.552041 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 138, prio 196, authtype none, intvl 1s, length 36
11:23:11.552046 IP 10.2.84.136 > 224.0.0.18: VRRPv2, Advertisement, vrid 137, prio 128, authtype none, intvl 1s, length 36

\$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
 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> mtu 1500 qdisc noop state DOWN
 link/void
5: eth0.84: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNOWN
 link/ether 00:18:51:ba:74:18 brd ff:ff:ff:ff:ff:ff
 inet 10.2.84.136/24 brd 10.2.84.255 scope global eth0.84
 inet6 fe80::218:51ff:feba:7418/64 scope link
 valid_lft forever preferred_lft forever
7: eth0.131: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNOWN
 link/ether 00:18:51:2d:bd:bf brd ff:ff:ff:ff:ff:ff
 inet xx.xx.xx.136/24 brd xx.xx.xx.255 scope global eth0.131
 inet6 fe80::218:51ff:fe2d:bdbf/64 scope link
 valid_lft forever preferred_lft forever

```
$ route -n
Kernel IP routing table
Destination     Gateway         Genmask        Flags Metric Ref  Use Iface
10.2.84.0       0.0.0.0        255.255.255.0  U     0      0      0 eth0.84
xx.xx.xx.0      0.0.0.0        255.255.255.0  U     0      0      0 eth0.131
192.168.0.0     10.2.84.1     255.255.0.0    UG    0      0      0 eth0.84
172.16.0.0      10.2.84.1     255.240.0.0    UG    0      0      0 eth0.84
10.0.0.0        10.2.84.1     255.0.0.0     UG    0      0      0 eth0.84
0.0.0.0          xx.xx.xx.1   0.0.0.0       UG    0      0      0 eth0.131
```

Subject: Re: Not seeing multicast on veth
Posted by [zajdee](#) on Tue, 20 Apr 2010 11:32:53 GMT
[View Forum Message](#) <=> [Reply to Message](#)

I see that noone has replied so far - I haven't found the solution either, anyways I am trying again and adding some more details:

Adding mcast address to veth interface manually does not work - is this a bug or a feature?

```
[host]:~ $ ip maddr add 224.0.0.18 dev veth2084066.0
ioctl: Invalid argument
[host]:~ $ ip maddr add 224.0.0.18 dev eth0.84

[host]:~ $ ip maddr show
<--cut-->
27: eth0.84
    link 33:33:ff:41:bb:22 users 2
    link 01:00:5e:00:00:01 users 2
    link 33:33:00:00:00:01 users 2
    inet 224.0.0.1
    inet6 ff02::1:ff41:bb22
    inet6 ff02::1
<--cut-->
34: veth2084066.0
    link 33:33:ff:4d:cd:f1
    link 01:00:5e:00:00:01
    link 33:33:00:00:00:01
    inet 224.0.0.1
    inet6 ff02::1:ff4d:cdf1
    inet6 ff02::1
```

Trying on an updated, custom built kernel:

```
[master]:~ $ uname -r
2.6.32-openvz
```

tcpdump on host:

```
13:29:44.661438 IP 10.2.84.66 > 224.0.0.22: igmp v3 report, 1 group record(s)
13:29:45.485942 IP 10.2.84.66.60990 > 10.2.34.28.123: NTPv4, Client, length 48
13:29:47.184397 IP 10.2.84.66 > 224.0.0.22: igmp v3 report, 1 group record(s)
13:30:18.202465 IP 10.2.84.66 > 224.0.0.22: igmp v3 report, 1 group record(s)
13:30:18.519178 IP 10.2.84.66.35496 > 10.2.34.3.123: NTPv4, Client, length 48
13:30:20.720397 IP 10.2.84.66 > 224.0.0.22: igmp v3 report, 1 group record(s)
```

tcpdump on guest:

```
13:29:44.661424 IP 10.2.84.66 > 224.0.0.22: igmp v3 report, 1 group record(s)
13:29:45.485942 IP 10.2.84.66.60990 > 10.2.34.28.123: NTPv4, Client, length 48
13:29:47.184397 IP 10.2.84.66 > 224.0.0.22: igmp v3 report, 1 group record(s)
13:30:00.384623 IP 10.2.84.66 > 224.0.0.18: VRRPv2, Advertisement, vrid 99, prio 0, authtype
none, intvl 5s, length 36
13:30:05.687839 IP 10.2.84.66 > 224.0.0.18: VRRPv2, Advertisement, vrid 99, prio 0, authtype
none, intvl 5s, length 36
13:30:11.159059 IP 10.2.84.66 > 224.0.0.18: VRRPv2, Advertisement, vrid 99, prio 0, authtype
none, intvl 5s, length 36
13:30:16.526877 IP 10.2.84.66 > 224.0.0.18: VRRPv2, Advertisement, vrid 99, prio 0, authtype
none, intvl 5s, length 36
13:30:18.202465 IP 10.2.84.66 > 224.0.0.22: igmp v3 report, 1 group record(s)
13:30:18.519178 IP 10.2.84.66.35496 > 10.2.34.3.123: NTPv4, Client, length 48
13:30:20.720397 IP 10.2.84.66 > 224.0.0.22: igmp v3 report, 1 group record(s)
```

As you can see, the igmp v3 report (join/leave) passes from guest to host (veth interface), however vrrp advertisements do not.
