
Subject: Re: Running DHCP on VPS, (on a router..)

Posted by [dev](#) on Fri, 28 Sep 2007 13:39:13 GMT

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

Dariusz Pietrzak wrote:

>>>HN: tcpdump -n -i eth0.107

>>

>>is it veth pair interface, right? ok...

>

> I might not exactly follow the term 'veth pair interface', but eth0.107

> and veth are joined together by a bridge:

by pair I mean that veth interface has 2 ends: one inside VE and another one in HN.

then what is eth0.107? vlan?

> br107 8000.0018511cf71d no eth0.107

> veth107.0

>

>

>>and is 192.168.8.254 assigned to veth inside VE?

>

> 192.168.8.251 is assigned to veth, .254 is a ucarp-managed IP and it's

> assigned to br107.

so, this means that your host system replies to DHCP request below.

cause this IP belongs to HN.

are you running DHCP server in HN as well?

>>>08:16:19.401880 00:1b:d4:7e:76:2a > 01:00:0c:cc:cc:cd SNAP Unnumbered, ui, Flags
[Command], length 50

>>>08:16:21.154240 IP 0.0.0.0.68 > 255.255.255.255.67: BOOTP/DHCP, Request from
00:1b:d5:2c:bf:38, length 308

>>>08:16:21.185096 IP 192.168.8.254.67 > 255.255.255.255.68: BOOTP/DHCP, Reply, length
300

>>>08:16:21.187344 arp who-has 192.168.9.254 tell 192.168.9.97

>>>

>>>

>>>HN: tcpdump -n -i br107

>>>08:16:19.401880 00:1b:d4:7e:76:2a > 01:00:0c:cc:cc:cd SNAP Unnumbered, ui, Flags
[Command], length 50

>>>08:16:21.154240 IP 0.0.0.0.68 > 255.255.255.255.67: BOOTP/DHCP, Request from
00:1b:d5:2c:bf:38, length 308

>>>08:16:21.185096 IP 192.168.8.254.67 > 255.255.255.255.68: BOOTP/DHCP, Reply, length
300

>>>08:16:21.187344 arp who-has 192.168.9.254 tell 192.168.9.97

>>>

>>>and finally, from inside the vps:

>>>VPS: tcpdump -n -i eth0

>>>08:16:19.401886 00:1b:d4:7e:76:2a > 01:00:0c:cc:cc:cd SNAP Unnumbered, ui, Flags
[Command], length 50
>>>08:16:21.185110 IP 192.168.8.254.67 > 255.255.255.255.68: BOOTP/DHCP, Reply, length
300
>>>08:16:21.187356 arp who-has 192.168.9.254 tell 192.168.9.97
>>>08:16:21.188600 arp who-has 192.168.8.1 tell 192.168.9.97
>>
>>I wonder why timestamp of Reply here is greater then timestamp in HN...
>
> yes, it looks strange, especially since the next packet - arp req, has
> similarly delayed timestamp, like something in veth delayed packets after
> dhcp req/reply sequence.
> This delay is visible here: .401880 on bridge and eth0.107, but .401886
> from veth. But this .000006 grows to .000014 and then shrinks to .000012
> after missing the packet.

if my logic above about DHCP server in HN is correct, then no wonder... :)

>>to DHCP requests and you see the reply only. Can you please check (if it is true - shutting
down
>>dhcp server inside VE won't affect tcpdump output)?
>
> This is true, but why do some packets visible on bridge and physical
> interface disappear on veth? When I disable the other dhcp server there is
> no reply at all. dhcpd doesn't see the packet.

I guess this can be due to veth filtering which I removed by the patch I send you.
This filtering drops the packets which should not be seen by VE.
If the patch helps - I'm right.

>>> as you can see, broadcast request from dhcp client is missing, but strangely enough, the
reply
>>>is clearly visible.
>>
>>Or it works, but only tcpdump is missing the packet?
>
> That would be possible, but still puzzling, and that wouldn't explaining why
> both tcpdump and dhcpd are missing the packet.
> Besides, this machine is very lightly loaded at the moment, and the
> traffic at hand is also abysmally small, not to mention that this happens quite regularly, almost
> all dhcp requests get lost and never arrive to my dhcpd.

Thanks,
Kirill
