

---

Subject: Re: VE IP address stops working after several hours

Posted by [fatbrother](#) on Thu, 18 Sep 2008 13:31:12 GMT

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

---

maratrus wrote on Thu, 18 September 2008 18:38Hi,

1. I'm not absolutely certain of my facts by I'm inclined to think that your HN never sends arp response even everything works great.

Other machines just have your VEs arp-record in their arp-table which occur there during the VE startup or changing ip address.

Could possibly check this fact by using tcpdump when everything works fine.

I though I've seen it responding to ARPs. But now I did clean check from another host and yes it does not respond, even when real\_ip2 is pingable.

After IP change it sends out "who has real\_ip2 tell real\_ip2", may be that's what Cisco router uses for creating it's APR record.

So I'm ready to accept that you're right.

BTW, I tried ip neigh del/add sequence you suggested and it does not help.

Quote:2. So, you have Real named table which contains only one record "default via real\_ip\_router dev eth0.425 table Real", don't you?

Yes. That's first line of my "ip route list table all":

```
>default via real_ip_router dev eth0.425 table Real
```

Quote:If yes, I don't understand how this configuration works at all.

If something comes to VE it comes to HN first then the packet should be routed to your VE but Real table doesn't contain any mention of your VE which is on the venet interface. But your rule says that everything that comes "from real\_ip\_network/28" should be passed through Real table.

Oops. Well, I can explain how it works. I never tried to access real\_ip2 from other hosts on real\_network. So all packets I've sent to real\_ip2 weren't "from real\_ip\_network/28".

Quote:3. If the Real table would contain the record like "real\_ip2 dev venet0 scope link" I can explain the first and the second points.

Could you possibly put the record like "real\_ip2 dev venet0 scope link" to the Real table and check if something change.

I added this record and the real\_ip2 become pingable from other hosts on real\_network!!

And it probably responded to ARP. I wasn't running tcpdump at that moment, but other hosts now have normal ARP records for real\_ip2. I need to wait for 4 hours to tell whether that solves main problem. I have no access to the router (at least not before our Cisco admin will come to the work tomorrow), so I cannot flush it's ARP cache and cannot do a clean check.

But I probably should add rule for "to real\_network" too, because now I get this:

```
>20:04:03.316468 IP real_ip3 > real_ip2: ICMP echo request, id 35397, seq 14, length 64
```

```
>20:04:03.316506 IP real_ip2 > real_ip3: ICMP echo reply, id 35397, seq 14, length 64
```

```
>20:04:03.316842 IP real_router > real_ip2: ICMP redirect real_ip3 to host real_ip3, length 36
```

and that redirect goes after every ping.

That's not a big issue, these hosts aren't supposed to communicate to each other over

real\_network, but that's annoying.

How do you suggest my ip rules should look like? I do not want to add a new explicit rule for every new real\_ip for my VE, but probably that's the only correct way...

Quote:4.

Quote:

32765: from real\_ip1 lookup Real

By the way why do you have this rule?

Err... I first added this rule, before I even added real\_ip2. I just was testing HN connectivity to VLAN. Then I added rule for the network, and I haven't rebooted HN since then.

---