
Subject: (openvz 2.6.18-53.1.4.el5.028stab053.4 + intel e1000) not sending arp notifies

Posted by [rickb](#) on Thu, 13 Mar 2008 18:43:55 GMT

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

Hi, I have one server running 2.6.18-53.1.4.el5.028stab053.4 64bit and using the ethernet card- "e1000: eth0: e1000_probe: Intel(R) PRO/1000 Network Connection".

The problem is that every few hours packets are not sent to the server, I believe because the nic is not broadcasting arp notifications on layer 2. I can use the 'arp send' to fix the situation by arping for every IP assigned to VEs, but its quite a bad way to fix the problem, and also seems to bring down networking in the VE when the arp send program runs.

I see the following in the VE syslog when running arp send:

```
Mar 13 02:25:02 vps network: Shutting down interface venet0: succeeded
```

```
Mar 13 02:25:03 vps network: Bringing up interface venet0: succeeded
```

Has anyone seen any situations similar to this and/or have any pointers on how I can ensure the server is sending arp broadcasts?

thanks!

Rick

Subject: Re: (openvz 2.6.18-53.1.4.el5.028stab053.4 + intel e1000) not sending arp notifies

Posted by [maratrus](#) on Fri, 14 Mar 2008 09:53:11 GMT

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

Hi,

just to clarify the situation:

1. You have got an Openvz server and several VE on it.
2. All VEs use venet0 interface.
3. Everything was OK but suddenly the following issue happened:
4. External computer tried to ping one of your VE and it was impossible.

Arp requests come from external machine to your server but your server don't reply because arp table doesn't contain appropriate record ("arp -n"). Am I right?

Then you try to use "arp send" utility from HN. Could you please specify what did you do? arp send -U -i VE_ID -c 1 eth0 or something? Or may be anything else? And after you've typed that command you could observe that messages in /var/log/messages inside VE.

You didn't restart your network interfaces on HN. But may be anything strange in logs on HN or in dmesg?

P.S. I'm terribly sorry for these questions but I really wasn't be able to clearly understand your

previous post.

Subject: Re: (openvz 2.6.18-53.1.4.el5.028stab053.4 + intel e1000) not sending arp notifies

Posted by [rickb](#) on Fri, 14 Mar 2008 14:22:19 GMT

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

Hi, thank you for the reply. I am happy to provide further information. The steps your provided are all correct. The arpsend command I use to bring the VEs back to life is:

```
/usr/sbin/arpsend -U -i <IP> -c 2 eth0
```

When I do this for a VE's IP, the VE can reach the network fine. Doing so does trigger the syslog message I included previously but that is not such a huge problem, just some information I added.

I never restart the network interface on the HN. My assumption is that the NIC is not dispatching arp packets to notify the switch and local machines of what IPs the hardware node is assigned often enough (or at all aside from the initial network script start) to keep the VEs online. It seems like the HN's IP is always online as I am always able to reach it, but ALL of the VE's IPs will become unreachable at the same time hours or a day later. I SSH into the HN and use the arpsend in a little for loop for all of the VE's IPs and everything is back to normal for a few hours.

What I did in the meantime was add a static arp in the switch to the HN's mac addresss, but it is a very bad solution because the IPs may be moved in the future.

Does the NIC send the arp packets automatically every XX minutes? What would cause this problem on a high level?, without thinking about openvz.

Thanks for any ideas!

Subject: Re: (openvz 2.6.18-53.1.4.el5.028stab053.4 + intel e1000) not sending arp notifies

Posted by [maratrus](#) on Mon, 17 Mar 2008 12:44:20 GMT

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

Hi,

Den wrote (http://forum.openvz.org/index.php?t=msg&goto=27782&# msg_27782)

Quote:The node will arp reply for 1.2.3.4 if and only if
ip r g 1.2.3.4 from [your_ip] dev [incoming dev]
will return a route _OTHER_ than one to [incoming dev]

So in our case HN have to reply if we have appropriate record in our route table. Please, check that you have it (I mean something like VE_IP dev venet0 scope link src HN_IP).

I have the following assumption:

in spite of the fact that we have that routing entries our proxy_arp variable (I mean /proc/sys/net/ipv4/conf/eth0/proxy_arp) is set to be 0. So HN doesn't response. But during the VE starting vzctl put appropriate proxy arp records and that is why HN can response to VE_IP arp requests. So we can add this record manually like vzctl utility does. (on HN)ip neigh add proxy VE_IP dev eth0

But who flush these records? I don't know.

Quote:Does the NIC send the arp packets automatically every XX minutes?
As far as I know NIC doesn't sent them.