

---

Subject: Run XORP in VE

Posted by [yoolee](#) on Tue, 19 Feb 2008 04:28:01 GMT

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

---

Hi, All

I would like to host multiple XORP instances using OpenVZ. But there is a strange problem: After xorp runs for 7-8 minutes I can not ping all interfaces in that particular VE, including lo interface. This is strange because xorp uses lo interface to communicate among its processes, it keeps running though I can not ping lo intf.

Xorp uses netlink socket to get interfaces information and update forwarding information in kernel. It also uses normal socket and raw socket to exchange routing messages.

I did the same thing using Xen before, and I would like to port previous work to OpenVZ for a large scale testing.

Could anyone give me some clues? Thanks a lot!

---

---

Subject: Re: Run XORP in VE

Posted by [den](#) on Tue, 19 Feb 2008 07:25:55 GMT

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

---

Have you checked link and routing state inside a VE when you see the problem?

Pls also say exact kernel version you are using.

Regards,  
Den

---

---

Subject: Re: Run XORP in VE

Posted by [yoolee](#) on Tue, 19 Feb 2008 15:30:15 GMT

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

---

Thanks for the reply.

I am using 2.6.18-fza-028stab051.1-686-bigmem, OS is debian 4.0.

When the problem occurs, routing table of that VE looks like:

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
192.0.2.1	*	255.255.255.255	UH	0	0	0	venet0
10.16.0.4	*	255.255.255.252	U	0	0	0	tun1
10.16.0.0	*	255.255.255.252	U	0	0	0	tun0
default	192.0.2.1	0.0.0.0	UG	0	0	0	venet0

And I can ping that VE from hardware node all the time, which is strange.

---

---

Subject: Re: Run XORP in VE  
Posted by [den](#) on Tue, 19 Feb 2008 15:43:04 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

could you pls run  
ip route list table local  
and  
ip a l  
ip l l

Local node can be pinged when appropriate record is present in the local routing table and lo is up and running even when you trying to ping any other interface.

Regards,  
Den

---

---

Subject: Re: Run XORP in VE  
Posted by [yoolee](#) on Tue, 19 Feb 2008 16:30:28 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Oh, those commands show nothing when the problem occurs.  
For some reason the routes are cleaned by xorp?

The following are the normal output of those commands:

```
# ip route list table local
broadcast 10.16.0.4 dev tun1 proto kernel scope link src 10.16.0.5
local 10.16.0.5 dev tun1 proto kernel scope host src 10.16.0.5
broadcast 127.255.255.255 dev lo proto kernel scope link src 127.0.0.1
broadcast 10.16.0.7 dev tun1 proto kernel scope link src 10.16.0.5
broadcast 10.16.0.0 dev tun0 proto kernel scope link src 10.16.0.2
local 10.10.0.102 dev venet0 proto kernel scope host src 10.10.0.102
local 10.16.0.2 dev tun0 proto kernel scope host src 10.16.0.2
broadcast 10.16.0.3 dev tun0 proto kernel scope link src 10.16.0.2
broadcast 127.0.0.0 dev lo proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo proto kernel scope host src 127.0.0.1
local 127.0.0.1 dev venet0 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo proto kernel scope host src 127.0.0.1
```

```
# ip l l
1: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue
```

```
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue
  link/void
13: tun0: <POINTOPOINT,MULTICAST,NOARP,UP,10000> mtu 1500 qdisc pfifo_fast qlen 500
  link/[65534]
15: tun1: <POINTOPOINT,MULTICAST,NOARP,UP,10000> mtu 1500 qdisc pfifo_fast qlen 500
  link/[65534]
```

```
# ip a l
```

```
1: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue
  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,UP,10000> mtu 1500 qdisc noqueue
  link/void
  inet 127.0.0.1/32 scope host venet0
  inet 10.10.0.102/32 scope global venet0:0
13: tun0: <POINTOPOINT,MULTICAST,NOARP,UP,10000> mtu 1500 qdisc pfifo_fast qlen 500
  link/[65534]
  inet 10.16.0.2/30 scope global tun0
15: tun1: <POINTOPOINT,MULTICAST,NOARP,UP,10000> mtu 1500 qdisc pfifo_fast qlen 500
  link/[65534]
  inet 10.16.0.5/30 scope global tun1
```

---

Subject: Re: Run XORP in VE  
Posted by [yoolee](#) on Tue, 19 Feb 2008 17:08:19 GMT  
[View Forum Message](#) <> [Reply to Message](#)

Well, it is not exact to say those commands show nothing.  
They just block there and can't return.

Because xorp also use netlink to get/set interfaces and routes information, it is probably a netlink problem.

---

Subject: Re: Run XORP in VE  
Posted by [den](#) on Wed, 20 Feb 2008 07:47:44 GMT  
[View Forum Message](#) <> [Reply to Message](#)

could you pls look into  
cat /proc/user\_beancounters  
you should see a faults there (last column).

So, you face a lack of resources for this VE.

Regards,  
Den

---

---

Subject: Re: Run XORP in VE  
Posted by [yoolee](#) on Wed, 20 Feb 2008 15:11:44 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Den, thank you very much.

Yes, there is a fault:

```
dgramrcvbuf      0  261696  262144  262144  38075
```

So I think I would increase that quota in my VE conf file, how large that parameter can be set?

The purpose of my experiment is to create at least 1,000 VEs for hosting XORP, and conduct large-scale routing test.

---

---

Subject: Re: Run XORP in VE  
Posted by [yoolee](#) on Wed, 20 Feb 2008 16:54:30 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I have one question now, from my observation the dgramrcvbuf parameter keeps increasing during my experiment runs, until it reach the quota. Why doesn't it get a chance to release?

My scenario is :

One VE acts as a openvpn client, HN is vpn server which connects to a remote machine.  
One xorp runs in that VE, only ospf is enabled in xorp conf. The behavior of ospf is simple, just sending hello packets in a multicast fashion.

---

---

Subject: Re: Run XORP in VE  
Posted by [den](#) on Thu, 21 Feb 2008 08:35:03 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

this sounds like a bug.

---

How many netlink sockets do you have?  
How many buffers are on them in reality?

Do you have an idea how to reproduce the problem using 'ip' without taking into account the XORP?

Regards,  
Den

---

---

Subject: Re: Run XORP in VE  
Posted by [yoollee](#) on Thu, 21 Feb 2008 21:27:03 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From my observation the parameter "dgramrcvbuf" depends on how many multicast packets xorp sends in that VE. For example if I enlarge the Hello-interval of OSPF protocol (the hello packet is a kind of multicast packet with dst 224.0.0.5) from 1 sec. to 2 sec., it takes "dgramrcvbuf" exactly twice as much time to reach the quota.

Now I set the limit of "dgramrcvbuf" to 680000, and I found "dgramrcvbuf" increases to about 620000 and stops increasing. So I can run my experiment for a long time.

Yes, I will try to find out an easy way to reproduce the problem for your debugging.

Thanks a lot!

---