
Subject: CT can'ts communicate beetwen them
Posted by Kernn on Mon, 31 Oct 2016 09:34:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi everybody,

I'm a new user of openvz and i need some help for understand what's happens on my server.

I have 4 CT on my server.

```
[root@myserver1 ~]# vzlist
  CTID   NPROC STATUS  IP_ADDR      CPUS HOSTNAME      NAME
  101    128 running x.x.x.1      8 xxxxxxxxxx   xxxxxx
  102    213 running x.x.x.2      4 xxxxxxxxxx   -
  103    58  running x.x.x.3      2 xxxxxxxxxx   xxxxxx
  104    42  running x.x.x.4      2 xxxxxxxxxx   xxxxx
```

All CT are connected by interface venet0.

IPs x.x.x.1, x.x.x.2, x.x.x.3, x.x.x.4 are on the same subnet.

After an update (yum update) and a reboot, i have observed that all CT can'ts communicate beetwen them. :(

For example :

On the global server, i can join 101 :

```
[root@myserver1 ~]# ping x.x.x.1
PING x.x.x.1 (x.x.x.1) 56(84) bytes of data:
64 bytes from x.x.x.1: icmp_seq=1 ttl=64 time=0.061 ms
64 bytes from x.x.x.1: icmp_seq=2 ttl=64 time=0.047 ms
64 bytes from x.x.x.1: icmp_seq=3 ttl=64 time=0.039 ms
^C
--- x.x.x.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2138ms
rtt min/avg/max/mdev = 0.039/0.049/0.061/0.009 ms
```

On the global server, i can join 102 :

```
[root@myserver1 ~]# ping x.x.x.2
PING x.x.x.2 (x.x.x.2) 56(84) bytes of data:
64 bytes from x.x.x.2: icmp_seq=1 ttl=64 time=0.040 ms
64 bytes from x.x.x.2: icmp_seq=2 ttl=64 time=0.052 ms
64 bytes from x.x.x.2: icmp_seq=3 ttl=64 time=0.046 ms
^C
--- x.x.x.2 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2189ms
```

rtt min/avg/max/mdev = 0.040/0.046/0.052/0.005 ms

But, in the CT 101, i can not join the CT 102

```
[root@myserver1 ~]# vzctl enter 101
entered into CT 101
```

```
[root@CT1 /]# ping x.x.x.2
PING x.x.x.2 (x.x.x.2) 56(84) bytes of data.
^C
--- x.x.x.2 ping statistics ---
3 packets transmitted, 0 received, 100% packet loss, time 2757ms
```

In the CT 102, i can not join the CT 101.

I have no REJECT iptables :

```
[root@myserver1 ~]# iptables -L -vn
Chain INPUT (policy ACCEPT 4812K packets, 418M bytes)
pkts bytes target prot opt in out source destination

Chain FORWARD (policy ACCEPT 36M packets, 22G bytes)
pkts bytes target prot opt in out source destination

Chain OUTPUT (policy ACCEPT 4557K packets, 8704M bytes)
pkts bytes target prot opt in out source destination
```

```
[root@CT1 /]# iptables -L -vn
Chain INPUT (policy ACCEPT 58218 packets, 9907K bytes)
pkts bytes target prot opt in out source destination
58218 9907K acctin all -- * * 0.0.0.0/0 0.0.0.0/0

Chain FORWARD (policy ACCEPT 0 packets, 0 bytes)
pkts bytes target prot opt in out source destination

Chain OUTPUT (policy ACCEPT 84621 packets, 92M bytes)
pkts bytes target prot opt in out source destination
84621 92M acctout all -- * * 0.0.0.0/0 0.0.0.0/0
```

```
Chain acctin (1 references)
pkts bytes target prot opt in out source destination
6728 2189K all -- * * 0.0.0.0/0 127.0.0.1
6728 2189K all -- * * 0.0.0.0/0 127.0.0.1
51490 7718K all -- * * 0.0.0.0/0 x.x.x.1
```

```
Chain acctout (1 references)
```

```

pkts bytes target  prot opt in  out  source      destination
6728 2189K      all -- *   *   127.0.0.1  0.0.0.0/0
6728 2189K      all -- *   *   127.0.0.1  0.0.0.0/0
77893 90M         all -- *   *   x.x.x.1    0.0.0.0/0

```

```

[root@CT2 /]# iptables -L -vn
Chain INPUT (policy ACCEPT 31M packets, 5475M bytes)
pkts bytes target  prot opt in  out  source      destination

Chain FORWARD (policy ACCEPT 0 packets, 0 bytes)
pkts bytes target  prot opt in  out  source      destination

Chain OUTPUT (policy ACCEPT 31M packets, 5821M bytes)
pkts bytes target  prot opt in  out  source      destination

```

On the global server i have this route :

```

[root@myserver1 ~]# route -n
Table de routage IP du noyau
Destination  Passerelle  Genmask      Indic Metric Ref  Use Iface
x.x.x.1     0.0.0.0     255.255.255.255 UH  0  0  0 venet0
x.x.x.2     0.0.0.0     255.255.255.255 UH  0  0  0 venet0
x.x.x.3     0.0.0.0     255.255.255.255 UH  0  0  0 venet0
x.x.x.4     0.0.0.0     255.255.255.255 UH  0  0  0 venet0
x.x.x.0     0.0.0.0     255.255.255.0  U  0  0  0 eth2
x.x.x.0     0.0.0.0     255.255.255.0  U  0  0  0 eth0
x.x.x.0     0.0.0.0     255.255.255.0  U  0  0  0 eth1
x.x.x.0     x.x.x.240   255.255.255.0  UG  0  0  0 eth1
x.x.x.0     x.x.x.254   255.255.255.0  UG  0  0  0 eth1
0.0.0.0     x.x.x.254   0.0.0.0        UG  0  0  0 eth0

```

```

[root@myserver1 ~]# ip rule show
0:    from all lookup local
10:   from x.x.x.253 blackhole
11:   from all lookup 99
12:   from x.x.x.0/24 iif venet0 lookup 181
13:   from all iif venet0 blackhole
32766: from all lookup main
32767: from all lookup default

```

In this case, this two CT (CT1 in x.x.x.1 and CT2 in x.x.x.2) are connected by venet0 and the route table seems good :

```

[root@CT2 /]# ip route get x.x.x.1
x.x.x.1 dev venet0 src x.x.x.2
      cache mtu 1500 hoplimit 64

```

```
[root@myserver1 ~]# ip route get x.x.x.1 from x.x.x.2
x.x.x.1 from x.x.x.2 dev venet0
  cache mtu 1500 hoplimit 64
```

```
[root@CT1 /]# ip route get x.x.x.1
local x.x.x.x dev lo src x.x.x.1
  cache <local> mtu 65520 hoplimit 64
```

With a tcpdump on "CT1", "myserver1" and "CT"2, i note that the packet forwarding seems blocked between myserver1 and CT1.

The global server "myserver1" receives packets on venet0 since CT2. He send packets of CT2 on venet0 and forward packet at destination of CT1 on venet0. But on venet0 of CT1, no packet arrive. :/

Somebody have an idee or a sugestion ?

Thanks you for your help.
