Subject: No route to container's private address.
Posted by slowpoke on Wed, 28 Mar 2018 21:22:47 GMT
View Forum Message <> Reply to Message

Hi guys.

I have an awkward situation here.
On my virtualization host I have multiple addresses assigned to a physical interface. These addresses are from different subnets. Network is managed by NetworkManager.
There's a bunch of containers running in host routed network mode with private addresses from 10.10.0.0/24 range.

The thing is, that all the CTs are pingable from Hypervizor by their private addresses but at the same time not all of them are actually accessible with let's say telnet or ssh.
When I try to telnet into such a container, I get "No route to host" error whilst another identical CT accepts the connection without any problems. And I'm sure that the service I'm trying to connect to, is up and running inside the CT.

Here's my hypervizor config (centos 7)

```
# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: enp1s0f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen
1000
    link/ether 00:25:90:c6:6f:06 brd ff:ff:ff:ff:ff:ff
    inet 123.123.81.239/24 brd 123.123.81.255 scope global enp1s0f0
       valid_lft forever preferred_lft forever
    inet 123.123.80.150/24 brd 123.123.80.255 scope global enp1s0f0
       valid_lft forever preferred_lft forever
    inet 123.123.81.142/24 brd 123.123.81.255 scope global secondary enp1s0f0
       valid_lft forever preferred_lft forever
    inet 123.123.80.16/24 brd 123.123.80.255 scope global secondary enp1s0f0
       valid_lft forever preferred_lft forever
    inet 123.123.81.146/24 brd 123.123.81.255 scope global secondary enp1s0f0
       valid_lft forever preferred_lft forever
    inet6 fe80::225:90ff:fec6:6f06/64 scope link
       valid_lft forever preferred_lft forever
3: enp1s0f1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen
1000
    link/ether 00:25:90:c6:6f:07 brd ff:ff:ff:ff:ff:ff
4: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
state UNKNOWN
```

```
    link/void
    inet6 fe80::1/128 scope link
       valid_lft forever preferred_lft forever
5: host-routed: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN qlen 1000
    link/ether ea:75:33:3b:24:3e brd ff:ff:ff:ff:ff:ff
6: virbr0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue state DOWN
qlen 1000
    link/ether 52:54:00:00:54:c6 brd ff:ff:ff:ff:ff:ff
    inet 10.37.130.2/24 brd 10.37.130.255 scope global virbr0
       valid_lft forever preferred_lft forever
    inet6 fdb2:2c26:f4e4::1/64 scope global
       valid_lft forever preferred_lft forever
    inet6 fe80::5054:ff:fe00:54c6/64 scope link
       valid_lft forever preferred_lft forever
7: virbr0-nic: <BROADCAST,MULTICAST> mtu 1500 qdisc mq master virbr0 state DOWN qlen
1000
    link/ether 52:54:00:00:54:c6 brd ff:ff:ff:ff:ff:ff
```

Routes:

```
# ip r
default via 123.123.81.1 dev enp1s0f0 proto static metric 100
10.10.0.2 dev venet0 scope link
10.10.0.3 dev venet0 scope link
10.10.0.4 dev venet0 scope link
10.10.0.5 dev venet0 scope link
10.10.0.6 dev venet0 scope link
10.37.130.0/24 dev virbr0 proto kernel scope link src 10.37.130.2
123.123.80.0/24 dev enp1s0f0 proto kernel scope link src 123.123.80.150 metric 100
123.123.81.0/24 dev enp1s0f0 proto kernel scope link src 123.123.81.239 metric 100
```

This happenes:

```
# ping 10.10.0.3
PING 10.10.0.3 (10.10.0.3) 56(84) bytes of data.
64 bytes from 10.10.0.3: icmp_seq=1 ttl=64 time=0.087 ms

# telnet 10.10.0.3 10500
Trying 10.10.0.3...
Connected to 10.10.0.3.

# ping 10.10.0.5
PING 10.10.0.5 (10.10.0.5) 56(84) bytes of data.
64 bytes from 10.10.0.5: icmp_seq=1 ttl=64 time=0.114 ms

# telnet 10.10.0.5 10500
```

Trying 10.10.0.5...
telnet: connect to address 10.10.0.5: No route to host

Please, help!

Subject: Re: No route to container's private address.
Posted by vaverin on Fri, 30 Mar 2018 07:47:49 GMT
View Forum Message <> Reply to Message

Seems you need to look at traffic on venet interface
tcpdump -i venet0 -nve

Do you probably have some firewall, on host or inside affected containers ?

Please show routing rules inside affected containers, probably they do not send proper reply back to host IP?

Btw I assume your containers have venet interface only.
Perhaps they have some other interfaces too ?

Thank you,
Vasily Averin
