

Fixing Docker and VPN IP Address Conflicts

by Nate Lampton // August 15, 2019

YOU MIGHT ALSO LIKE

Creating Default Docker Networks

(<https://github.com/docker/compose/issues/4336>).

Occasionally when accessing a client network, I encounter a situation where certain servers are not accessible despite everyone else on the team being able to access the same domain, either by HTTP or SSH. It turns out a frequent culprit of this problem is Docker and its networking mechanisms.

When you start up Docker, it appropriates some IP addresses for its own usage. These are usually in the local networking space, which include the following:

10.0.0.0 to 10.255.255.255

172.16.0.0 to 172.31.255.255

192.168.0.0 to 192.168.255.255

Discovering Conflicts

You can easily find which IP addresses are being used by your local network or VPN by using `route -n`. Connect to your VPN *first* then run this command.

```

1 $ route -n
2
3 Kernel IP routing table
4 Destination      Gateway            Genmask           Flags Metric Ref    Use
5 Iface
6 0.0.0.0           192.168.1.1       0.0.0.0           UG        600    0      0
7 wlp2s0
8 128.2.5.132      192.168.1.1       255.255.255.255  UGH       600    0      0
9 wlp2s0
10 172.16.0.0        0.0.0.0           255.255.0.0       U         50     0      0
11 vpn0
12 172.18.6.0        0.0.0.0           255.255.255.0     U         50     0      0
13 vpn0
14 172.18.11.0       0.0.0.0           255.255.255.0     U         50     0      0
15 vpn0
16 172.19.0.0        0.0.0.0           255.255.0.0       U         50     0      0
17 vpn0
18 172.19.238.0      0.0.0.0           255.255.255.0     U         50     0      0
19 vpn0
20 172.19.247.0      0.0.0.0           255.255.255.0     U         50     0      0
21 vpn0
22 172.19.249.0      0.0.0.0           255.255.255.0     U         50     0      0
23 vpn0
24 172.20.0.0        0.0.0.0           255.255.0.0       U         50     0      0
25 vpn0
26 172.20.40.0       0.0.0.0           255.255.254.0     U         50     0      0
27 vpn0
   172.20.42.0       0.0.0.0           255.255.254.0     U         50     0      0
   vpn0
   172.20.45.0       0.0.0.0           255.255.255.0     U         50     0      0
   vpn0
   172.20.46.0       0.0.0.0           255.255.255.0     U         50     0      0
   vpn0

```

172.21.0.0	0.0.0.0	255.255.0.0	U	50	0	0
vpn0						
172.22.8.0	0.0.0.0	255.255.255.0	U	50	0	0
vpn0						
172.22.25.0	0.0.0.0	255.255.255.0	U	50	0	0
vpn0						
172.22.114.0	0.0.0.0	255.255.255.0	U	50	0	0
vpn0						
172.22.115.0	0.0.0.0	255.255.255.0	U	50	0	0
vpn0						
172.24.2.0	0.0.0.0	255.255.255.0	U	50	0	0
vpn0						
172.24.238.0	0.0.0.0	255.255.255.0	U	50	0	0
vpn0						
172.29.48.0	0.0.0.0	255.255.240.0	U	50	0	0
vpn0						
172.29.80.0	0.0.0.0	255.255.240.0	U	50	0	0
vpn0						

Looking at this above data, we can derive that the IP address between the ranges of 172.16.x.x and 172.29.x.x are not safe for docker to use.

To determine what IP addresses docker itself is using, we can use the `ip addr` command to see what addresses the networking bridges claim.

Prior to starting docker and any containers:

```

1  $ ip addr
2
3  1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group
4  default qlen 1000
5      link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
6      inet 127.0.0.1/8 scope host lo
7          valid_lft forever preferred_lft forever
8      inet6 ::1/128 scope host
9          valid_lft forever preferred_lft forever
10 2: wlp2s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state
11  UP group default qlen 1000
12      link/ether 9c:b6:d0:92:d5:b1 brd ff:ff:ff:ff:ff:ff
13      inet 192.168.1.245/24 brd 192.168.1.255 scope global dynamic
14  noprefixroute wlp2s0
15          valid_lft 85363sec preferred_lft 85363sec
16      inet6 fe80::69fe:7762:4ecd:d5ae/64 scope link noprefixroute
17          valid_lft forever preferred_lft forever
18 3: enx8cae4cf13353: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc
19  fq_codel state DOWN group default qlen 1000
20      link/ether 8c:ae:4c:f1:33:53 brd ff:ff:ff:ff:ff:ff
21 4: vpn0: <POINTOPOINT,MULTICAST,NOARP,UP,LOWER_UP> mtu 1300 qdisc fq_codel
22  state UP group default qlen 500
      link/none
      inet 172.31.235.27/16 brd 172.31.255.255 scope global noprefixroute
      vpn0
          valid_lft forever preferred_lft forever
      inet6 fe80::5b9c:a58c:5b5a:6b90/64 scope link stable-privacy
          valid_lft forever preferred_lft forever

```

And then after starting docker:

```
1 $ ip addr
2
3 ip addr
4 1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group
5 default qlen 1000
6     link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
7     inet 127.0.0.1/8 scope host lo
8         valid_lft forever preferred_lft forever
9     inet6 ::1/128 scope host
10        valid_lft forever preferred_lft forever
11 2: wlp2s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state
12 UP group default qlen 1000
13     link/ether 9c:b6:d0:92:d5:b1 brd ff:ff:ff:ff:ff:ff
14     inet 192.168.1.245/24 brd 192.168.1.255 scope global dynamic
15 noprefixroute wlp2s0
16         valid_lft 85977sec preferred_lft 85977sec
17     inet6 fe80::69fe:7762:4ecd:d5ae/64 scope link noprefixroute
18         valid_lft forever preferred_lft forever
19 3: enx8cae4cf13353: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc
20 fq_codel state DOWN group default qlen 1000
21     link/ether 8c:ae:4c:f1:33:53 brd ff:ff:ff:ff:ff:ff
22 4: vpn0: <POINTOPOINT,MULTICAST,NOARP,UP,LOWER_UP> mtu 1300 qdisc fq_codel
23 state UP group default qlen 500
24     link/none
25     inet 172.31.235.27/16 brd 172.31.255.255 scope global noprefixroute
26 vpn0
27         valid_lft forever preferred_lft forever
28     inet6 fe80::5b9c:a58c:5b5a:6b90/64 scope link stable-privacy
29         valid_lft forever preferred_lft forever
30 5: br-05743ccfd659: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc
31 noqueue state DOWN group default
32     link/ether 02:42:d8:4d:41:60 brd ff:ff:ff:ff:ff:ff
33     inet 172.20.0.1/16 brd 172.20.255.255 scope global br-05743ccfd659
```

```
34         valid_lft forever preferred_lft forever
35 6: br-08e37aab0021: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc
36 noqueue state DOWN group default
37     link/ether 02:42:40:a3:a0:63 brd ff:ff:ff:ff:ff:ff
38     inet 172.21.0.1/16 brd 172.21.255.255 scope global br-08e37aab0021
39         valid_lft forever preferred_lft forever
40 8: br-70d04f7b2a8c: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc
41 noqueue state DOWN group default
42     link/ether 02:42:35:bb:bc:eb brd ff:ff:ff:ff:ff:ff
43     inet 172.19.0.1/16 brd 172.19.255.255 scope global br-70d04f7b2a8c
        valid_lft forever preferred_lft forever
9: br-86768d2533cf: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc
noqueue state DOWN group default
    link/ether 02:42:71:49:8e:d7 brd ff:ff:ff:ff:ff:ff
    inet 172.18.0.1/16 brd 172.18.255.255 scope global br-86768d2533cf
        valid_lft forever preferred_lft forever
10: docker0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
state DOWN group default
    link/ether 02:42:a0:61:83:8f brd ff:ff:ff:ff:ff:ff
    inet 172.240.0.1/24 brd 172.240.0.255 scope global docker0
        valid_lft forever preferred_lft forever
```

We can see from the above that Docker has claimed 172.18.0.1/16 - 172.21.0.1/16. Which will make it so that routes specified by the VPN or local network will not work.

How to fix this?

First list your current docker networks.

```
$ docker network list
NETWORK ID NAME DRIVER SCOPE
28881c0a72ad cmu_default bridge local
b736a4c00275 host host local
e87ba6af1530 lando_bridge_network bridge local
7d9a9e0a3797 landoproxyhyperion5000gandalfedition_edge bridge
local
4601ca10be74 none null local
```

Those "bridge" entries map to the "br-" prefixed entries seen in the `ip addr` listing.

You can determine which networks are claiming which subnets by using `docker network inspect 28881c0a72ad`. In the JSON output you'll see something similar to this:

1	"Subnet": "172.18.0.0/24",
2	"Gateway": "172.18.0.1"

If you have set up these bridges yourself simply removing the networks and recreating them with a different subnet may be sufficient:

1	<code>docker network rm 05743ccfd659</code>
2	<code>docker network create --driver=bridge --subnet=192.168.0.0/16 br0</code>

Better yet, it's a good idea to prevent docker from automatically claiming subnets that are in conflict with your local network. To fix this you may modify the default daemon.json file used by docker.

```
1  $ vi /etc/docker/daemon.json
2  {
3      "default-address-pools" : [
4          {
5              "base" : "172.240.0.0/16",
6              "size" : 24
7          }
8      ]
9  }
```

Then newly created networks will automatically fit within the specified "base" and "size" parameters.

In my situation, I was using Lando, a docker management tool. So after removing each network and updating the configuration daemon.json configuration file, I had it regenerate all the networks from scratch with:

```
1  $ lando rebuild
```

Now all my networks are free from conflicts.

Note that individual projects can already specify different subnets in their docker-compose.yml or .lando.yml files. However, with different users needing different networks, control at that level may not be desirable.

More information about setting the default networks can be found in this Docker compose issue:

<https://github.com/docker/compose/issues/4336> (<https://github.com/docker/compose/issues/4336>). Happy networking!

Nate Lam ton

Nate Lampton is a leader in Open Source with over a decade of contributions to the Drupal project. He joined Lullabot in 2006.

© 2020 Lullabot, Inc.