
Subject: tcp packet loss on loopback

Posted by [chernomor](#) on Wed, 03 Feb 2016 09:18:24 GMT

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

1. kernel 2.6.32-042stab112.15
2. zero failcnt in /proc/user_beancounters
3. no iptables rules

ip a

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
2: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
state UNKNOWN
    link/void
    inet 127.0.0.1/32 scope host venet0
    inet 10.10.1.57/32 brd 10.10.1.57 scope global venet0:0
```

ping 127.0.0.1

```
PING 127.0.0.1 (127.0.0.1) 56(84) bytes of data.
64 bytes from 127.0.0.1: icmp_seq=1 ttl=64 time=0.016 ms
^C
```

--- 127.0.0.1 ping statistics ---

```
1 packets transmitted, 1 received, 0% packet loss, time 880ms
rtt min/avg/max/mdev = 0.016/0.016/0.016/0.000 ms
```

Ping is OK.

Test case #1

telnet 127.0.0.1 4534

Trying 127.0.0.1...

telnet: connect to address 127.0.0.1: Connection refused

It's OK - no anybody listen port 4534

Test case #2

strace nc -l 127.0.0.1 4534

....

```
socket(PF_INET, SOCK_STREAM, IPPROTO_TCP) = 3
setsockopt(3, SOL_SOCKET, SO_REUSEADDR, [1], 4) = 0
setsockopt(3, SOL_SOCKET, SO_REUSEPORT, [1], 4) = 0
bind(3, {sa_family=AF_INET, sin_port=htons(4534), sin_addr=inet_addr("127.0.0.1")}, 16) = 0
listen(3, 1)                                = 0
accept(3,
```

(in other session)

telnet 127.0.0.1 4534

Trying 127.0.0.1...

telnet: connect to address 127.0.0.1: Connection timed out

(in other session)

12:27:16.380522 IP localhost.localdomain.61060 > localhost.localdomain.4534: Flags [S], seq 2170016745, win 65483, options [mss 65495,sackOK,TS val 1048118367 ecr

1218379606,nop,wscale 8], length 0

12:27:17.379498 IP localhost.localdomain.61060 > localhost.localdomain.4534: Flags [S], seq 2170016745, win 65483, options [mss 65495,sackOK,TS val 1048119367 ecr

1218379606,nop,wscale 8], length 0

12:27:19.379492 IP localhost.localdomain.61060 > localhost.localdomain.4534: Flags [S], seq 2170016745, win 65483, options [mss 65495,sackOK,TS val 1048121367 ecr

1218379606,nop,wscale 8], length 0

It's NOT OK!

nc listen port 4534, but accept() sleep and not return!

I run tcpdump and see packets from telnet to nc, but not from nc to telnet.

Container restart does not solve problem.
