
Subject: [SOLVED] IPTables connection tracking stopped working.

Posted by [non7top](#) on Sat, 19 Oct 2013 12:49:58 GMT

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

Hello everyone.

Just about yesterday I decided to reboot several of my containers just to realize that iptables --state rules stopped working. I still have one container running without reboot(for 23 days) which has these rules working as expected. I don't think I made changes on containers iptables rules, but did make some changes on VE0, including several os updates. Ping from inside the container work just fine.

Here are the relevant config outputs.

VE0 iptables/route

iptables-save

Generated by iptables-save v1.4.7 on Sat Oct 19 16:38:33 2013

*mangle

:PREROUTING ACCEPT [104195:74172610]

:INPUT ACCEPT [22158:2977388]

:FORWARD ACCEPT [82037:71195222]

:OUTPUT ACCEPT [17925:31501333]

:POSTROUTING ACCEPT [99962:102696555]

COMMIT

Completed on Sat Oct 19 16:38:33 2013

Generated by iptables-save v1.4.7 on Sat Oct 19 16:38:33 2013

*filter

:INPUT ACCEPT [0:0]

:FORWARD ACCEPT [0:0]

:OUTPUT ACCEPT [17925:31501333]

-A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT

-A INPUT -p icmp -j ACCEPT

-A INPUT -i lo -j ACCEPT

-A INPUT -p tcp -m state --state NEW -m tcp --dport 45510 -j ACCEPT

-A INPUT -p tcp -m state --state NEW -m tcp --dport 80 -j ACCEPT

-A INPUT -p tcp -m state --state NEW -m tcp --dport 8080:8081 -j ACCEPT

-A INPUT -p tcp -m state --state NEW -m tcp --dport 20:21 -j ACCEPT

-A INPUT -p tcp -m state --state NEW -m tcp --dport 30000:31000 -j ACCEPT

-A INPUT -p udp -m udp --dport 33434:33523 -j REJECT --reject-with icmp-port-unreachable

-A INPUT -p udp -m udp --dport 44450:44500 -j REJECT --reject-with icmp-port-unreachable

-A INPUT -j REJECT --reject-with icmp-host-prohibited

-A FORWARD -s 37.200.71.8/29 -j ACCEPT

-A FORWARD -d 37.200.71.8/29 -j ACCEPT

-A FORWARD -s 192.168.11.0/24 -j ACCEPT

-A FORWARD -d 192.168.11.0/24 -j ACCEPT

-A FORWARD -j REJECT --reject-with icmp-host-prohibited

COMMIT

Completed on Sat Oct 19 16:38:33 2013

Generated by iptables-save v1.4.7 on Sat Oct 19 16:38:33 2013

```
*nat
:PREROUTING ACCEPT [2912:166427]
:POSTROUTING ACCEPT [1639:93157]
:OUTPUT ACCEPT [98:5998]
-A PREROUTING -d 37.200.71.10/32 -i eth0 -p tcp -m tcp --dport 10080 -j DNAT --to-destination
192.168.11.1:80
-A PREROUTING -d 37.200.71.10/32 -i eth0 -p tcp -m tcp --dport 45511 -j DNAT --to-destination
192.168.11.1:22
-A PREROUTING -d 37.200.71.10/32 -i eth0 -p tcp -m tcp --dport 10082 -j DNAT --to-destination
192.168.11.2:80
-A PREROUTING -d 37.200.71.10/32 -i eth0 -p tcp -m tcp --dport 45512 -j DNAT --to-destination
192.168.11.2:22
-A POSTROUTING -s 192.168.11.1/32 -o eth0 -j SNAT --to-source 37.200.71.10
-A POSTROUTING -s 192.168.11.2/32 -o eth0 -j SNAT --to-source 37.200.71.10
COMMIT
# Completed on Sat Oct 19 16:38:33 2013
```

```
# ip r
192.168.11.1 dev venet0 scope link
37.200.71.14 dev venet0 scope link
37.200.71.13 dev venet0 scope link
37.200.71.12 dev venet0 scope link
37.200.71.11 dev venet0 scope link
37.200.71.8/29 dev eth0 proto kernel scope link src 37.200.71.10
default via 37.200.71.9 dev eth0
```

```
# lsmod|grep -e ip_ -e xt_ -e nf_ -e ipt_
xt_length      1338  0
xt_hl          1547  0
xt_tcpmss      1623  0
xt_TCPMSS      3461  0
iptables_mangle 3493  0
xt_multiport   2716  0
xt_limit       2134  0
xt_dscp        2073  0
nf_nat_ftp     3523  0
nf_conntrack_ftp 12929 1 nf_nat_ftp
xt_DSCP        2849  0
xt_state       1508  10
iptables_filter 2937  5
iptables_nat   6302  1
nf_nat         23213 3 vzrst,nf_nat_ftp,iptables_nat
nf_conntrack_ipv4 9946 13 iptables_nat,nf_nat
nf_conntrack    80524 8 vzrst,vzcpd,nf_nat_ftp,nf_conntrack_ftp,xt_state,iptables_nat
,nf_nat,nf_conntrack_ipv4
nf_defrag_ipv4  1531  1 nf_conntrack_ipv4
ip_tables      18151 3 iptables_mangle,iptables_filter,iptables_nat
```

```

vz.conf
# cat /etc/vz/vz.conf
## Global parameters
VIRTUOZZO=yes
LOCKDIR=/vz/lock
DUMPPDIR=/vz/dump
VE0CPUUNITS=1000
VE_STOP_MODE=suspend

## Logging parameters
LOGGING=yes
LOGFILE=/var/log/vzctl.log
LOG_LEVEL=0
VERBOSE=0

## Disk quota parameters
DISK_QUOTA=yes
VZFASTBOOT=no

# Disable module loading. If set, vz initscript does not load any modules.
#MODULES_DISABLED=yes

# The name of the device whose IP address will be used as source IP for CT.
# By default automatically assigned.
#VE_ROUTE_SRC_DEV="eth0"

# Uncomment to limit CT IP ARP announces only to network interfaces
# having IPs within the same IP network as a container IP.
# Leave commented out to use all interfaces.
#NEIGHBOUR_DEVS=detect

## Fail if there is another machine in the network with the same IP
ERROR_ON_ARPFAIL="no"

## Template parameters
TEMPLATE=/vz/template

## Defaults for containers
VE_ROOT=/vz/root/$VEID
VE_PRIVATE=/vz/private/$VEID
CONFIGFILE="vswap-256m"
DEF_OSTEMPLATE="centos-6-x86"
NAMESERVER=inherit # Copy from host system's /etc/resolv.conf
## Filesystem layout for new CTs: either simfs (default) or ploop
#VE_LAYOUT=ploop

```

```
# User namespace configuration
LOCAL_UID=100000
LOCAL_GID=100000

## Load vzwdog module
VZWDOG="no"

## IPv4 iptables kernel modules to be enabled in CTs by default
IPTABLES="ipt_REJECT ipt_tos ipt_limit ipt_multiport iptable_filter iptable_mangle ipt_TCPMSS
ipt_tcpmss ipt_ttl ipt_length"
## IPv4 iptables kernel modules to be loaded by init.d/vz script
IPTABLES_MODULES="$IPTABLES"

## Enable IPv6
IPV6="no"

## IPv6 ip6tables kernel modules
IP6TABLES="ip6_tables ip6table_filter ip6table_mangle ip6t_REJECT"
```

Working VE101
iptables

```
# iptables -vnL
Chain INPUT (policy DROP 0 packets, 0 bytes)
pkts bytes target prot opt in out source destination
19M 43G ACCEPT all -- lo * 0.0.0.0/0 0.0.0.0/0
0 0 ACCEPT all -- venet0:0 * 0.0.0.0/0 0.0.0.0/0
58M 15G ACCEPT all -- venet0 * 0.0.0.0/0 0.0.0.0/0 state
RELATED,ESTABLISHED
200 7124 REJECT udp -- * * 0.0.0.0/0 0.0.0.0/0 udp dpts:33434:33523
reject-with icmp-port-unreachable
0 0 REJECT udp -- * * 0.0.0.0/0 0.0.0.0/0 udp dpts:44450:44500
reject-with icmp-port-unreachable
1563K 89M ALLOW all -- * * 0.0.0.0/0 0.0.0.0/0
19503 1292K LOGDROP 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 103M packets, 143G bytes)
pkts bytes target prot opt in out source destination

Chain ALLOW (1 references)
pkts bytes target prot opt in out source destination
95 5668 ACCEPT tcp -- venet0 * 0.0.0.0/0 0.0.0.0/0 tcp dpt:45510
1476K 83M ACCEPT tcp -- venet0 * 0.0.0.0/0 0.0.0.0/0 tcp dpt:80
84 4000 ACCEPT tcp -- venet0 * 0.0.0.0/0 0.0.0.0/0 tcp dpt:8080
```

```

21689 1210K ACCEPT  tcp -- venet0 *    0.0.0.0/0      0.0.0.0/0      tcp dpts:20:21
45216 3588K ACCEPT  icmp -- venet0 *    0.0.0.0/0      0.0.0.0/0

```

Chain LOGDROP (1 references)

```

pkts bytes target  prot opt in  out  source      destination
19503 1292K LOG      all -- *   *    0.0.0.0/0   0.0.0.0/0      LOG flags 0 level 7 prefix
`IPT_DROP: '
19503 1292K DROP     all -- *   *    0.0.0.0/0   0.0.0.0/0

```

Non working VE102

iptables

iptables -vnL

Chain INPUT (policy DROP 0 packets, 0 bytes)

```

pkts bytes target  prot opt in  out  source      destination
8222 9208K ACCEPT  all -- lo   *    0.0.0.0/0   0.0.0.0/0
0 0 ACCEPT  all -- venet0:0 *    0.0.0.0/0   0.0.0.0/0
0 0 ACCEPT  all -- *   *    0.0.0.0/0   0.0.0.0/0      state
RELATED,ESTABLISHED
0 0 REJECT  udp -- *   *    0.0.0.0/0   0.0.0.0/0      udp dpts:33434:33523
reject-with icmp-port-unreachable
0 0 REJECT  udp -- *   *    0.0.0.0/0   0.0.0.0/0      udp dpts:44450:44500
reject-with icmp-port-unreachable
11603 817K ALLOW   all -- *   *    0.0.0.0/0   0.0.0.0/0
68 10972 LOGDROP  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 37542 packets, 48M bytes)

```

pkts bytes target  prot opt in  out  source      destination

```

Chain ALLOW (1 references)

```

pkts bytes target  prot opt in  out  source      destination
103 10396 ACCEPT  tcp -- venet0 *    0.0.0.0/0      0.0.0.0/0      tcp dpt:45510
11391 792K ACCEPT  tcp -- venet0 *    0.0.0.0/0      0.0.0.0/0      tcp dpt:80
0 0 ACCEPT  tcp -- venet0 *    0.0.0.0/0      0.0.0.0/0      tcp dpt:8080
0 0 ACCEPT  tcp -- venet0 *    0.0.0.0/0      0.0.0.0/0      tcp dpts:20:21
0 0 ACCEPT  tcp -- venet0 *    0.0.0.0/0      0.0.0.0/0      tcp dpt:53
39 2718 ACCEPT  udp -- venet0 *    0.0.0.0/0      0.0.0.0/0      udp dpt:53
2 148 ACCEPT  icmp -- venet0 *    0.0.0.0/0      0.0.0.0/0

```

Chain LOGDROP (1 references)

```

pkts bytes target  prot opt in  out  source      destination
68 10972 LOG      all -- *   *    0.0.0.0/0   0.0.0.0/0      LOG flags 0 level 7 prefix
`IPT_DROP: '

```

68 10972 DROP all -- * * 0.0.0.0/0 0.0.0.0/0

From dmesg, blocked by LOGDROP

[5618.085257] IPT_DROP: IN=venet0 OUT= MAC= SRC=8.8.8.8 DST=37.200.71.12 LEN=72
TOS=0x00 PREC=0x00 TTL=51 ID=55483 PROTO=UDP SPT=53 DPT=46732 LEN=52
[5618.092405] IPT_DROP: IN=venet0 OUT= MAC= SRC=8.8.8.8 DST=37.200.71.12 LEN=121
TOS=0x00 PREC=0x00 TTL=51 ID=56508 PROTO=UDP SPT=53 DPT=46732 LEN=101
[5619.105489] IPT_DROP: IN=venet0 OUT= MAC= SRC=8.8.8.8 DST=37.200.71.12 LEN=179
TOS=0x00 PREC=0x00 TTL=51 ID=55484 PROTO=UDP SPT=53 DPT=50782 LEN=159

Subject: Re: IPTables connection tracking stopped working.

Posted by [non7top](#) on Sun, 20 Oct 2013 13:19:27 GMT

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

Nevermind. It was the vz.conf missing the nf_conntrack module. I had to restart the vz service first and then restart the vm itself for the setting to get applied.
