
Subject: Host firewall -- SOLVED

Posted by [ferp2](#) on Mon, 08 May 2006 15:44:08 GMT

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

Hello,

I have a generic firewall script that normally runs fine, but is giving me access problems between the host system and the vps. The firewall is set up for connections tracking using:

/sbin/modprobe ip_conntrack "ip_conntrack_enable_ve0=1"

With the host system firewall activated:

- I cannot ping a vps from the host system. The message I get is:

ping: sendmsg: Operation not permitted

- I cannot ping the host system from the vps either.
- I can successfully ping a separate machine on the same subnet from the host system.
- I can successfully ping the host system from a separate machine on the same subnet.

With the INPUT and OUTPUT policy set to ACCEPT, or the host system firewall deactivated:

- I can successfully ping a vps from the host system.
- I can successfully ping the host system from the vps.

How do I maintain the INPUT and OUTPUT policy set to DROP on the host system, and yet still be able to successfully ping the vps from the host system and vice-versa.

Thank you.

Subject: Re: Host firewall

Posted by [dev](#) on Tue, 09 May 2006 15:30:00 GMT

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

1. you are running 2.6.8 kernel, correct?
 2. can you show your rules in host system please, with iptables -L?
 3. or even get an access to your node, so I could quickly resolve it?
-

Subject: Re: Host firewall

Posted by [ferp2](#) on Tue, 09 May 2006 15:52:07 GMT

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

Here's the info you requested.

```
uname -r  
2.6.8-022stab064-up
```

```
/sbin/iptables -L -n  
Chain INPUT (policy DROP)  
target  prot opt source          destination  
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0      state RELATED,ESTABLISHED  
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0  
ACCEPT  icmp --  192.168.0.0/24  192.168.0.7    icmp type 8  
ACCEPT  tcp  --  192.168.0.2    192.168.0.7    tcp dpt:22 state NEW  
LOG     icmp --  0.0.0.0/0      192.168.0.7    icmp !type 8 LOG flags 0 level 4  
LOG     tcp  --  0.0.0.0/0      192.168.0.7    LOG flags 0 level 4
```

```
Chain FORWARD (policy ACCEPT)  
target  prot opt source          destination
```

```
Chain OUTPUT (policy DROP)  
target  prot opt source          destination  
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0      state RELATED,ESTABLISHED  
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0  
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0  
ACCEPT  icmp --  192.168.0.7    192.168.0.0/24  icmp type 0  
LOG     all  --  0.0.0.0/0      0.0.0.0/0      LOG flags 0 level 4
```

Thanks

Subject: Re: Host firewall
Posted by [dev](#) **on** Tue, 09 May 2006 18:14:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

are these rules:
ACCEPT all -- 0.0.0.0/0 0.0.0.0/0
intentional?

1. can you also please specify your network settings: host/VPS IP/mask?
 2. Is it possible to login to VPS via ssh from host (when policy is DROP and ping doesn't work)?
-

Subject: Re: Host firewall
Posted by [ferp2](#) **on** Tue, 09 May 2006 19:56:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

- 1.

are these rules:

ACCEPT all -- 0.0.0.0/0 0.0.0.0/0

intentional?

If I'm not mistaken the above comes from:

\$IPT -I INPUT 1 -p ALL -m state --state RELATED,ESTABLISHED -j ACCEPT

\$IPT -I OUTPUT 1 -p ALL -m state --state RELATED,ESTABLISHED -j ACCEPT

above 2 rules allow response to future rules using --state NEW

2.

can you also please specify your network settings: host/VPS IP/mask?

Host IP and mask:

ifconfig eth0

```
eth0      Link encap:Ethernet HWaddr 00:50:FC:72:56:44
          inet addr:192.168.0.7 Bcast:192.168.0.255 Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
          RX packets:14857 errors:0 dropped:0 overruns:0 frame:0
          TX packets:10341 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1562177 (1.4 MiB) TX bytes:1615613 (1.5 MiB)
          Interrupt:11 Base address:0xdc00
```

VPS IP and mask

ifconfig -a

```
lo      Link encap:Local Loopback
          inet addr:127.0.0.1 Mask:255.0.0.0
          UP LOOPBACK RUNNING MTU:16436 Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

```
venet0      Link encap:UNSPEC HWaddr 00-00-FF-00-00-00-00-00-00-00-00-00-00-00-00-00
          inet addr:127.0.0.1 P-t-P:127.0.0.1 Bcast:0.0.0.0 Mask:255.255.255.255
          UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
          RX packets:12798 errors:0 dropped:0 overruns:0 frame:0
          TX packets:8817 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:1182031 (1.1 MiB) TX bytes:1226952 (1.1 MiB)
```

```
venet0:0      Link encap:UNSPEC HWaddr 00-00-FF-FF-FF-FF-00-00-00-00-00-00-00-00-00-00-00
          inet addr:192.168.0.102 P-t-P:192.168.0.102 Bcast:0.0.0.0 Mask:255.255.255.255
          UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

Subject: Re: Host firewall

Posted by [Vasily Tarasov](#) on Wed, 10 May 2006 09:10:49 GMT

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

ferp2, unfortunately I was not able to reproduce your problem locally:

```
[root@dhcp0-82 ~]# uname -a
Linux dhcp0-82.sw.ru 2.6.8-022stab064.1 #1 Thu Jan 19 22:16:02 MSK 2006 i686 i686 i386
GNU/Linux
```

```
[root@dhcp0-82 ~]# iptables -L -n
Chain FORWARD (policy ACCEPT)
target  prot opt source          destination
target  prot opt source          destination
Chain INPUT (policy DROP)
target  prot opt source          destination
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0      state RELATED,ESTABLISHED
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0
ACCEPT  icmp --  192.168.0.0/24  192.168.0.82    icmp type 8
ACCEPT  tcp  --  192.168.0.254  192.168.0.82    tcp dpt:22 state NEW
LOG    icmp --  0.0.0.0/0      192.168.0.82    icmp !type 8 LOG flags 0 level 4
LOG    tcp  --  0.0.0.0/0      192.168.0.82    LOG flags 0 level 4

Chain OUTPUT (policy DROP)
target  prot opt source          destination
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0      state RELATED,ESTABLISHED
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0
ACCEPT  all  --  0.0.0.0/0      0.0.0.0/0
ACCEPT  icmp --  192.168.0.82  192.168.0.0/24  icmp type 0
LOG    all  --  0.0.0.0/0      0.0.0.0/0      LOG flags 0 level 4
```

```
[root@dhcp0-82 ~]# cat /etc/sysconfig/iptables
*filter
:INPUT DROP [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT DROP [0:0]
-A INPUT -s 0.0.0.0/0 -d 0.0.0.0/0 -m state --state RELATED,ESTABLISHED -j ACCEPT
-A INPUT -s 0.0.0.0/0 -d 0.0.0.0/0 -j ACCEPT
-A INPUT -p icmp --icmp-type 8 -s 192.168.0.0/24 -d 192.168.0.82 -j ACCEPT
-A INPUT -p tcp --destination-port 22 -s 192.168.0.254 -d 192.168.0.82 -m state --state NEW -j ACCEPT
-A INPUT -p icmp --icmp-type ! 8 -s 0.0.0.0/0 -d 192.168.0.82 -j LOG --log-level 4
-A INPUT -p tcp -s 0.0.0.0/0 -d 192.168.0.82 -j LOG --log-level 4

-A OUTPUT -s 0.0.0.0/0 -d 0.0.0.0/0 -m state --state RELATED,ESTABLISHED -j ACCEPT
-A OUTPUT -s 0.0.0.0/0 -d 0.0.0.0/0 -j ACCEPT
-A OUTPUT -s 0.0.0.0/0 -d 0.0.0.0/0 -j ACCEPT
```

```
-A OUTPUT -p icmp --icmp-type 0 -s 192.168.0.82 -d 192.168.0.0/24 -j ACCEPT
-A OUTPUT -s 0.0.0.0/0 -d 0.0.0.0/0 -j LOG --log-level 4
COMMIT
```

```
[root@dhcp0-82 ~]# ifconfig
eth0    Link encap:Ethernet HWaddr 00:0C:29:2A:0C:8A
        inet addr:192.168.0.82 Bcast:192.168.3.255 Mask:255.255.252.0
              UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
              RX packets:7752 errors:0 dropped:0 overruns:0 frame:0
              TX packets:502 errors:0 dropped:0 overruns:0 carrier:0
              collisions:0 txqueuelen:1000
              RX bytes:823851 (804.5 KiB) TX bytes:63597 (62.1 KiB)
              Interrupt:18 Base address:0x1080

lo     Link encap:Local Loopback
        inet addr:127.0.0.1 Mask:255.0.0.0
              UP LOOPBACK RUNNING MTU:16436 Metric:1
              RX packets:8 errors:0 dropped:0 overruns:0 frame:0
              TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
              collisions:0 txqueuelen:0
              RX bytes:560 (560.0 b) TX bytes:560 (560.0 b)

venet0   Link encap:UNSPEC HWaddr FF-BF-78-F6-BF-F0-AC-00-00-00-00-00-00-00-00-00
         UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
         RX packets:83 errors:0 dropped:0 overruns:0 frame:0
         TX packets:11 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:5628 (5.4 KiB) TX bytes:924 (924.0 b)
```

in VPS:

```
-bash-3.00# ifconfig
lo     Link encap:Local Loopback
        inet addr:127.0.0.1 Mask:255.0.0.0
              UP LOOPBACK RUNNING MTU:16436 Metric:1
              RX packets:8 errors:0 dropped:0 overruns:0 frame:0
              TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
              collisions:0 txqueuelen:0
              RX bytes:560 (560.0 b) TX bytes:560 (560.0 b)

venet0   Link encap:UNSPEC HWaddr FF-BF-38-F7-FF-BF-38-93-00-00-00-00-00-00-00-00-00
         inet addr:127.0.0.1 P-t-P:127.0.0.1 Bcast:0.0.0.0 Mask:255.255.255.255
         UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
         RX packets:7 errors:0 dropped:0 overruns:0 frame:0
         TX packets:43 errors:0 dropped:0 overruns:0 carrier:0
         collisions:0 txqueuelen:0
         RX bytes:588 (588.0 b) TX bytes:2940 (2.8 KiB)
```

```
venet0:0 Link encap:UNSPEC HWaddr FF-BF-38-F7-FF-BF-AC-F7-00-00-00-00-00-00-00-00  
inet addr:192.168.0.142 P-t-P:192.168.0.142 Bcast:192.168.0.142 Mask:255.255.255.255  
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

Can you, please, check iptables settings inside VPS?

May be there is some firewall inside VPS?

Thanks.

Subject: Re: Host firewall

Posted by [ferp2](#) on Wed, 10 May 2006 13:03:18 GMT

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

vps:

Chain INPUT (policy ACCEPT)
target prot opt source destination

Chain FORWARD (policy ACCEPT)
target prot opt source destination

Chain OUTPUT (policy ACCEPT)
target prot opt source destination

Tried the following:

set host input policy to accept and output policy to deny = vps can ping host, host cannot ping vps.

set host output policy to accept and input policy to deny = vps cannot ping host, host can ping vps.

Here are the host iptables rules:

```
# Path to executable  
IPT="/sbin/iptables"
```

```
# Enable OVZ kernel conntracks in host system  
/sbin/modprobe ip_conntrack "ip_conntrack_enable_ve0=1"
```

```
# Enable FTP connection tracking  
#/sbin/modprobe ip_conntrack_ftp
```

```
# Open ports for limited access  
OPENPORTS="22"
```

```
# INTERFACES  
INTERFACE="eth0" # Internet-connected interface  
LOOPBACK_INTERFACE="lo" # Loopback interface
```

```

IPADDR="192.168.0.7"

# NETWORKS
LOOPBACK="127.0.0.0/8"          # reserved loopback address range
CLASS_A="10.0.0.0/8"            # class A private networks
CLASS_B="172.16.0.0/12"          # class B private networks
CLASS_C="192.168.0.0/16"         # class C private networks
CLASS_D_MULTICAST="224.0.0.0/4"    # class D multicast addresses
CLASS_E_RESERVED_NET="240.0.0.0/5"  # class E reserved addresses
BROADCAST_SRC="0.0.0.0"           # broadcast source address
BROADCAST_DEST="255.255.255.255"   # broadcast destination address

# SUBNET
LAN="192.168.0.0/24"

# PORTS
PRIVPORTS="0:1023"              # privileged port range
UNPRIVPORTS="1024:65535"         # unprivileged port range

# =====
# Reset chains and set policies
# =====

# Remove any existing rules from all chains
$IPT -t filter --flush
$IPT -t nat --flush
$IPT -t mangle --flush

# Set default policy for all chains
# filter
$IPT --policy INPUT DROP
$IPT --policy OUTPUT DROP
$IPT --policy FORWARD ACCEPT

# Don't set nat and mangle tables to DROP unless
# you know what you're doing
# nat
#$$IPT -t nat --policy PREROUTING DROP
#$$IPT -t nat --policy OUTPUT DROP
#$$IPT -t nat --policy POSTROUTING DROP

# mangle
#$$IPT -t mangle --policy PREROUTING DROP
#$$IPT -t mangle --policy OUTPUT DROP

# Remove any pre-existing user-defined chains
$IPT -t filter --delete-chain
#$$IPT -t nat --delete-chain

```

```

#\$IPT -t mangle --delete-chain

# =====
# Using connection state to by-pass rule checking
# =====

# Using the state module alone, INVALID will break protocols that use
# bi-directional connections or multiple connections or exchanges,
# unless an ALG is provided for the protocol. At this time, FTP and
# IRC are the only protocols with ALG support.

\$IPT -I INPUT 1 -p ALL -m state --state RELATED,ESTABLISHED -j ACCEPT
\$IPT -I OUTPUT 1 -p ALL -m state --state RELATED,ESTABLISHED -j ACCEPT
#\$IPT -A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
#\$IPT -A OUTPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
# above 2 rules allow response to future rules using --state NEW

# Give this computer unrestricted access to the internet
\$IPT -A OUTPUT -p ALL -o \$INTERFACE -j ACCEPT

# Set traffic on the loopback interface to unrestricted
\$IPT -A INPUT -i \$LOOPBACK_INTERFACE -j ACCEPT
\$IPT -A OUTPUT -o \$LOOPBACK_INTERFACE -j ACCEPT

# =====
# Allow lan to ping host
#\$IPT -A INPUT -i \$INTERFACE -p icmp \
#--icmp-type echo-request -s \$LAN \
#-d \$IPADDR -m state --state NEW -j ACCEPT

# Allow LAN/PORTA to ping host
\$IPT -A INPUT -i \$INTERFACE -p icmp -s \$LAN \
--icmp-type echo-request -d \$IPADDR -j ACCEPT

\$IPT -A OUTPUT -o \$INTERFACE -p icmp -s \$IPADDR \
--icmp-type echo-reply -d \$LAN -j ACCEPT
# =====

# Allow limited access to host
for f in \$OPENPORTS; do
    \$IPT -A INPUT -i \$INTERFACE -p tcp \
    -s \$LAN -d \$IPADDR --dport \$f \
    -m state --state NEW -j ACCEPT
done

```

Hope this helps.

Subject: Re: Host firewall

Posted by [Vasily Tarasov](#) on Thu, 11 May 2006 08:42:26 GMT

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

Hello, the reason of this issue is that VPS and Host talk through interface venet0. So you have to tune iptables to work properly with this interface too.

Here is a small patch for your script.

You can easily change it if you wish some slightly other behaviour.

```
--- firewall.sh.back 2006-05-10 19:23:09.000000000 +0400
+++ firewall.sh 2006-05-10 20:40:00.000000000 +0400
@@ -97,6 +97,16 @@ $IPT -A INPUT -i $INTERFACE -p icmp -s $
$IPT -A OUTPUT -o $INTERFACE -p icmp -s $IPADDR \
--icmp-type echo-reply -d $LAN -j ACCEPT
+
+# Allow host to ping VPS
+VENET="venet0"
+VPSIP="192.168.0.102"
+$IPT -A OUTPUT -o $VENET -p icmp -s $IPADDR \
--icmp-type echo-request -d $VPSIP -j ACCEPT
+# Allow host to receive ping from VPS
+$IPT -A INPUT -i $VENET -p icmp -s $VPSIP \
--icmp-type echo-request -d $IPADDR -j ACCEPT
+
# =====
# Allow limited access to host
```

Good luck.

Subject: Re: Host firewall

Posted by [ferp2](#) on Thu, 11 May 2006 12:44:43 GMT

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

Yes, that's the solution. I should have realized communication between the host and vps occurs on the venet0 interface.

Thanks for your help.
