Subject: Networking/IPTables, cannot ping domains names from container with iptables on in HZ
Posted by openxs on Thu, 31 Jul 2008 19:11:00 GMT
View Forum Message <> Reply to Message

IPTables problem 31/07/08

I am running a server from home using dynDNS. I can ping internal/external IP addresses from my VPS (which is called 101) without a problem, but if I try to ping a domain mane it will not resolve and I get: unknown host google.co.uk.

Sorry if this question has already been answered, I have found posts with similar problems but the set up is generally different. I'm pretty sure I have not configured the iptables correctly.

I come to this conclusion because If I switch off IPtables on the HN then restart VPS 101, domains start to resolve on VPS 101, but I'm not sure that this is a good way to run the server...

This is what I have set up:

uname -rm = 2.6.18-53.1.19.el5.028stab053.14ent i686
vzctl version 3.0.22
HN = CentOS 5 with IP: 192.168.1.2
VPS 101 = CentOS 5 with IP: 192.168.1.5
Router = 192.168.1.1 (I have reserved 192.168.1.2 - 19 for static addresses)

I set this up following the quick start guide on the wiki, but I was a little uncertain about /etc/sysctl.conf, I have added the contents of my file below.

I also tried this from the OpenVZ wiki. Ref:
http://wiki.openvz.org/Using_NAT_for_container_with_private_ IPs

** How to provide access for container to Internet **

To enable the containers, which have only internal IP addresses, to access the Internet, SNAT (Source Network Address Translation, also known as IP masquerading) should be configured on the Hardware Node. This is ensured by the standard Linux iptables utility. To perform a simple SNAT setup, execute the following command on the Hardware Node:

iptables -t nat -A POSTROUTING -s src_net -o eth0 -j SNAT --to ip_address

Mine looks like this:

iptables -t nat -A POSTROUTING -s 192.168.1.5/19 -o eth0 -j SNAT --to 192.168.1.2

I have turned Iptables off so i can carry on using just my hardware firewall, do I actually need IPTables on the HZ? I would feel happier using/learning it. Am I missing something, I have to admit I have never really had to play with IPTables before so this is unchartered territory for me.

I found this post in the forums, but these guys solved the problem by switching IPTables off...
Ref: http://forum.openvz.org/index.php?t=msg&goto=11896&

Here are the contents of the files I modified during the install.

---

cat /etc/modprobe.conf

```
options ip_conntrack ip_conntrack_enable_ve0=1
alias eth0 tg3
alias scsi_hostadapter ata_piix
```

---

cat /etc/sysctl.conf

```
# Kernel sysctl configuration file for Red Hat Linux
#
# For binary values, 0 is disabled, 1 is enabled.  See sysctl( and
# sysctl.conf(5) for more details.

# Controls IP packet forwarding
net.ipv4.ip_forward = 1
net.ipv4.conf.default.proxy_arp = 0

# Controls source route verification
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.all.rp_filter = 1

# Do not accept source routing
net.ipv4.conf.default.accept_source_route = 0

# Controls the System Request debugging functionality of the kernel
kernel.sysrq = 1

# Controls whether core dumps will append the PID to the core filename
# Useful for debugging multi-threaded applications
kernel.core_uses_pid = 1

# Controls the use of TCP syncookies
net.ipv4.tcp_syncookies = 1

# Controls the maximum size of a message, in bytes
kernel.msgmnb = 65536

# Controls the default maxmimum size of a mesage queue
kernel.msgmax = 65536

# Controls the maximum shared segment size, in bytes
kernel.shmmax = 4294967295
```

```
# Controls the maximum number of shared memory segments, in pages
kernel.shmall = 268435456

# We do not want all our interfaces to send redirects
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
```

Here is some other information that might be useful:

Commands run on the HZ:

```
# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:21:5A:51:39:75
          inet addr:192.168.1.2  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::221:5aff:fe51:3975/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:443 errors:0 dropped:0 overruns:0 frame:0
          TX packets:333 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:39090 (38.1 KiB)  TX bytes:51661 (50.4 KiB)
          Interrupt:177

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          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 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
          UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
          RX packets:4 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:268 (268.0 b)  TX bytes:380 (380.0 b)

# ip route list table all

192.168.1.5 dev venet0  scope link
192.168.1.0/24 dev eth0  proto kernel  scope link  src 192.168.1.2
169.254.0.0/16 dev eth0  scope link
default via 192.168.1.1 dev eth0
broadcast 192.168.1.0 dev eth0  table 255  proto kernel  scope link  src 192.168.1.2
```

```
broadcast 127.255.255.255 dev lo  table 255  proto kernel  scope link  src 127.0.0.1
local 192.168.1.2 dev eth0  table 255  proto kernel  scope host  src 192.168.1.2
broadcast 192.168.1.255 dev eth0  table 255  proto kernel  scope link  src 192.168.1.2
broadcast 127.0.0.0 dev lo  table 255  proto kernel  scope link  src 127.0.0.1
local 127.0.0.1 dev lo  table 255  proto kernel  scope host  src 127.0.0.1
local 127.0.0.0/8 dev lo  table 255  proto kernel  scope host  src 127.0.0.1
fe80::/64 dev eth0  metric 256  expires 21334181sec mtu 1500 advmss 1440 hoplimit
4294967295
unreachable default dev lo  table unspec  proto none  metric -1  error -101 hoplimit 255
local ::1 via :: dev lo  table 255  proto none  metric 0  mtu 16436 advmss 16376 hoplimit
4294967295
local fe80::221:5aff:fe51:3975 via :: dev lo  table 255  proto none  metric 0  mtu 16436 advmss
16376 hoplimit 4294967295
ff00::/8 dev eth0  table 255  metric 256  expires 21334181sec mtu 1500 advmss 1440 hoplimit
4294967295
unreachable default dev lo  table unspec  proto none  metric -1  error -101 hoplimit 255

#  iptables -t nat -L && iptables -t filter -L && iptables -t mangle -L

Chain PREROUTING (policy ACCEPT)
target     prot opt source               destination

Chain POSTROUTING (policy ACCEPT)
target     prot opt source               destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source               destination
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
RH-Firewall-1-INPUT  all  --  anywhere           anywhere

Chain FORWARD (policy ACCEPT)
target     prot opt source               destination
RH-Firewall-1-INPUT  all  --  anywhere           anywhere

Chain OUTPUT (policy ACCEPT)
target     prot opt source               destination

Chain RH-Firewall-1-INPUT (2 references)
target     prot opt source               destination
ACCEPT     all  --  anywhere             anywhere
ACCEPT     icmp --  anywhere             anywhere            icmp any
ACCEPT     esp  --  anywhere             anywhere
ACCEPT     ah   --  anywhere             anywhere


ACCEPT     udp  --  anywhere             224.0.0.251         udp dpt:mdns
ACCEPT     udp  --  anywhere             anywhere            udp dpt:ipp
```

```
ACCEPT     tcp  --  anywhere            anywhere           tcp dpt:ipp
ACCEPT     all  --  anywhere            anywhere           state RELATED,ESTABLISHED
ACCEPT     tcp  --  anywhere            anywhere           state NEW tcp dpt:ssh
ACCEPT     tcp  --  anywhere            anywhere           state NEW tcp dpt:http
ACCEPT     tcp  --  anywhere            anywhere           state NEW tcp dpt:https
REJECT     all  --  anywhere            anywhere           reject-with icmp-host-prohibited
Chain PREROUTING (policy ACCEPT)
target     prot opt source               destination

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

Chain POSTROUTING (policy ACCEPT)
target     prot opt source               destination


# arp -n
Address                  HWtype  HWaddress           Flags Mask            Iface
192.168.1.21             ether   00:19:7E:21:74:82   C                     eth0
192.168.1.1              ether   00:18:F8:4B:6D:96   C                     eth0
192.168.1.5              *       *                   MP                    eth0

# ip a l
2: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
4: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:21:5a:51:39:75 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.2/24 brd 192.168.1.255 scope global eth0
    inet6 fe80::221:5aff:fe51:3975/64 scope link
       valid_lft forever preferred_lft forever
6: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/void

# ip r l
192.168.1.5 dev venet0  scope link
192.168.1.0/24 dev eth0  proto kernel  scope link  src 192.168.1.2
169.254.0.0/16 dev eth0  scope link
```

default via 192.168.1.1 dev eth0

Now the same commands on the VPS 101...

```
# ifconfig
lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          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 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
          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:4 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:380 (380.0 b)  TX bytes:268 (268.0 b)

venet0:0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
          inet addr:192.168.1.5  P-t-P:192.168.1.5  Bcast:192.168.1.5  Mask:255.255.255.255
          UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1

# ip route list table all
192.0.2.0/24 dev venet0  scope host
169.254.0.0/16 dev venet0  scope link
default via 192.0.2.1 dev venet0
broadcast 127.255.255.255 dev lo  table 255  proto kernel  scope link  src 127.0.0.1
local 192.168.1.5 dev venet0  table 255  proto kernel  scope host  src 192.168.1.5
broadcast 192.168.1.5 dev venet0  table 255  proto kernel  scope link  src 192.168.1.5
broadcast 127.0.0.0 dev lo  table 255  proto kernel  scope link  src 127.0.0.1
local 127.0.0.1 dev lo  table 255  proto kernel  scope host  src 127.0.0.1
local 127.0.0.1 dev venet0  table 255  proto kernel  scope host  src 127.0.0.1
local 127.0.0.0/8 dev lo  table 255  proto kernel  scope host  src 127.0.0.1
unreachable default dev lo  table unspec  proto none  metric -1  error -101 hoplimit 255
local ::1 via :: dev lo  table 255  proto none  metric 0  mtu 16436 advmss 16376 hoplimit
4294967295
unreachable default dev lo  table unspec  proto none  metric -1  error -101 hoplimit 255

# iptables -t nat -L && iptables -t filter -L && iptables -t mangle -L
iptables v1.3.5: can't initialize iptables table `nat': Table does not exist (do you need to insmod?)
```

Perhaps iptables or your kernel needs to be upgraded.

```
# ip a l
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/void
    inet 127.0.0.1/32 scope host venet0
    inet 192.168.1.5/32 brd 192.168.1.5 scope global venet0:0

# ip r l
192.0.2.0/24 dev venet0  scope host
169.254.0.0/16 dev venet0  scope link
default via 192.0.2.1 dev venet0
```

If there is more information I can provide that may help then let me know, thanks.

Subject: Re: Networking/IPTables, cannot ping domains names from container with iptables on in HZ
Posted by kir on Fri, 01 Aug 2008 12:12:21 GMT
View Forum Message <> Reply to Message

Do you have correct nameserver set in /etc/resolv.conf inside a VE?

Subject: Re: Networking/IPTables, cannot ping domains names from container with iptables on in HZ
Posted by openxs on Fri, 01 Aug 2008 14:06:51 GMT
View Forum Message <> Reply to Message

Hi, thanks for the reply. Yes, I believe the nameserver is OK, I am using my router as the name server.

Running cat on the VE.

```
# cat /etc/resolv.conf
nameserver 192.168.1.1
```

This is the same as HZ's resolve.conf too. I have tried it using external nameservers, i.e. openDNS's, but this didn't help. Once IP tables are off, everything resolves fine.

I have set up a webserver on the VE with Ajaxterm running, everything is working as expected...

any other idea's?

I assume that you looked at my post and couldn't see anything obviously wrong there?

Subject: Re: Networking/IPTables, cannot ping domains names from container with iptables on in HZ
Posted by kir on Fri, 01 Aug 2008 14:31:19 GMT
View Forum Message <> Reply to Message

You have a rule for IPP (Internet Printing Protocol) which works over UDP:
ACCEPT udp -- anywhere anywhere udp dpt:ipp

DNS name resolving also works over UDP, so you need to add a similar rule, only use 'dpt:domain' at the end. This should help.

Subject: Re: Networking/IPTables, cannot ping domains names from container with iptables on in HZ
Posted by openxs on Mon, 01 Dec 2008 11:26:50 GMT
View Forum Message <> Reply to Message

My apologies, I should have replied to this post ages ago, it worked and resolved my problem at the time, so I have been using OpenVZ for 6 months with no trouble. However, on a reinstall I have the same problem.

I added the rule again, but it still don't seem to work, have I put it in the wrong place?

```
iptables -L
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
RH-Firewall-1-INPUT  all  --  anywhere             anywhere

Chain FORWARD (policy ACCEPT)
target     prot opt source               destination
RH-Firewall-1-INPUT  all  --  anywhere             anywhere

Chain OUTPUT (policy ACCEPT)
target     prot opt source               destination

Chain RH-Firewall-1-INPUT (2 references)
target     prot opt source               destination
ACCEPT     all  --  anywhere             anywhere
ACCEPT     icmp --  anywhere             anywhere            icmp any
ACCEPT     esp  --  anywhere             anywhere
ACCEPT     ah   --  anywhere             anywhere
ACCEPT     udp  --  anywhere             224.0.0.251         udp dpt:mdns
```

```
ACCEPT     udp  --  anywhere             anywhere          udp dpt:ipp
ACCEPT     tcp  --  anywhere             anywhere          tcp dpt:ipp
ACCEPT     all  --  anywhere             anywhere          state RELATED,ESTABLISHED
ACCEPT     tcp  --  anywhere             anywhere          state NEW tcp dpt:ssh
REJECT     all  --  anywhere             anywhere          reject-with icmp-host-prohibited
ACCEPT     tcp  --  anywhere             anywhere          tcp dpt:domain
ACCEPT     udp  --  anywhere             anywhere          udp dpt:domain
```

Subject: Re: Networking/IPTables, cannot ping domains names from container with iptables on in HZ
Posted by hostzilla on Sun, 28 Mar 2010 01:56:08 GMT
View Forum Message <> Reply to Message

iptables -I RH-Firewall-1-INPUT -p udp -m udp --dport 53 -j ACCEPT

Subject: Re: Networking/IPTables, cannot ping domains names from container with iptables on in HZ
Posted by hostzilla on Tue, 30 Mar 2010 22:26:43 GMT
View Forum Message <> Reply to Message

the command required is:
iptables -I RH-Firewall-1-INPUT -p udp -m udp --dport 53 -j ACCEPT