Subject: Iptables question
Posted by jvgrago on Tue, 28 Mar 2006 01:06:32 GMT
View Forum Message <> Reply to Message

I have just installed OpenVZ and the install was going well until I had to run vzpkgcache -f and this is when I noticed that with the OpenVZ kernel I have no access outside this server. If I stop iptables, I can do a host google.com and get a response or even do the vzpkgcache -f and it works just fine. At the top of my iptables, it shows this:

vi /etc/sysconfig/iptables

Generated by iptables-save v1.3.0 on Mon Mar 27 19:42:46 2006
*nat
:OUTPUT ACCEPT [0:0]
:POSTROUTING ACCEPT [0:0]
:PREROUTING ACCEPT [0:0]
COMMIT
Completed on Mon Mar 27 19:42:46 2006
Generated by iptables-save v1.3.0 on Mon Mar 27 19:42:46 2006
*mangle
:FORWARD ACCEPT [0:0]
:INPUT ACCEPT [277:27444]
:OUTPUT ACCEPT [226:23424]
:POSTROUTING ACCEPT [226:23424]
:PREROUTING ACCEPT [277:27444]
COMMIT
Completed on Mon Mar 27 19:42:46 2006
Generated by iptables-save v1.3.0 on Mon Mar 27 19:42:46 2006
*filter
:FORWARD ACCEPT [0:0]
:INPUT ACCEPT [0:0]
:OUTPUT ACCEPT [226:23424]
:RH-Firewall-1-INPUT - [0:0]
-A FORWARD -j RH-Firewall-1-INPUT

I dont see anything wrong with this, but as soon as I start iptables, I lose internet on this server.

Second issue:

Another issue that I have ran into is that the virtual ethernet venet0 does not activate on boot. If I try to activate it, It does not let me. Here is an output of an ifconfig.

eth0 Link encap:Ethernet HWaddr 00:0C:41:E8:AD:22
inet addr:192.168.0.9 Bcast:192.168.0.255 Mask:255.255.255.0
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:34800 errors:0 dropped:0 overruns:0 frame:0
TX packets:33813 errors:0 dropped:0 overruns:0 carrier:0

collisions:0 txqueuelen:1000
RX bytes:33516583 (31.9 MiB) TX bytes:5089116 (4.8 MiB)
Interrupt:11 Base address:0xdc00

eth1 Link encap:Ethernet HWaddr 00:0F:B5:08:9A:2D
inet addr:192.168.0.99 Bcast:192.168.0.255 Mask:255.255.255.0
UP BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:24 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b) TX bytes:2280 (2.2 KiB)
Interrupt:10 Base address:0xe000

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:55 errors:0 dropped:0 overruns:0 frame:0
TX packets:55 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:3946 (3.8 KiB) TX bytes:3946 (3.8 KiB)

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:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:7 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

I have deactivated eth1, but for some reason its still pulling an ip (with no cable attached either).

Any ideas?

Thanks,
Jim

Subject: Re: Iptables question
Posted by dev on Tue, 28 Mar 2006 23:55:23 GMT
View Forum Message <> Reply to Message

1. which kernel do you use? if 2.6.8 have you tried this from the FAQ:
Quote:
Q. My node is unaccessible through the network after reboot...

A. You need to check your firewall rules. The problem is that default stateful firewall rules are not available on the host system. To make this functionality available, load the ip_conntrack module with the additional parameter "ip_conntrack_enable_ve0=1". However, this method is highly not

recommended because tracking all the connection on the host system lead to performance degradation, more memory usage and also may lead to the total server inaccessibility due to reaching of the overall connection limit.

?

2. sorry, I didn't get your question/explanation about venet0. As far as I can see from your ifconfig output eth1 is activated, and I can't catch how is it related to venet0...
can you please make it more clear? what messages are printed when it fails to activate? why do you think is it not activated (in your output it is UP!) ?. Which commands do you use to activate it and which fail?

Subject: Re: Iptables question
Posted by jvgrago on Wed, 29 Mar 2006 20:36:32 GMT
View Forum Message <> Reply to Message

Hi, Eth1 was being activated at boot time and was manually assigned an ip, I wasnt sure if it effect anything else thats why I mentioned it. I am using the 2.6.8 Kernel (Latest for fedora core 4 that you have). I rebooted the server and watched to see what messages are being displayed relating to Venet0. It tells me this:

Device Venet0 does not seem to be present, Delaying initialization.

Now, an ifconfig shows this:

```
eth0      Link encap:Ethernet  HWaddr 00:0C:41:E8:AD:22
          inet addr:192.168.0.9  Bcast:192.168.0.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:61 errors:0 dropped:0 overruns:0 frame:0
          TX packets:82 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:11099 (10.8 KiB)  TX bytes:12268 (11.9 KiB)
          Interrupt:11 Base address:0xdc00

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:14 errors:0 dropped:0 overruns:0 frame:0
          TX packets:14 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:888 (888.0 b)  TX bytes:888 (888.0 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:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:6 overruns:0 carrier:0
```

collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)

Should Venet0 have an ip? Can I manually assign one if so?
I tried the ip_conntrack method you suggested, and still no outside activity if iptables is running.
When I restart networking, I see the following:

Shutting down interface eth0: [OK]
Shutting down interface venet0: [OK]
Shutting down loopback interface: [OK]
Disabling IPv4 packet forwarding: net.ipv4.ip_forward = 0
[OK]
Bringing up loopback interface: [OK]
Bringing up interface eth0: [OK]
Bringing up interface eth1:
Determining IP information for eth1... failed; no link present. Check cable?
[FAILED]
Bringing up interface venet0:
Determining IP information for venet0... failed.
[FAILED]

Any suggestions?

Thanks,
Jim

Subject: Re: Iptables question
Posted by dev on Thu, 30 Mar 2006 13:54:44 GMT
View Forum Message <> Reply to Message

Maybe this is the answer:
Quote:
Disabling IPv4 packet forwarding: net.ipv4.ip_forward = 0

you have to set ip forwarding to 1 in /etc/sysctl.conf
See installation guide.

Subject: Re: Iptables question
Posted by jvgrago on Thu, 30 Mar 2006 14:53:37 GMT
View Forum Message <> Reply to Message

My sysctl.conf file did contain the IP forwarding = 1. Here is a copy of it.

```
# 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
# 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
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
```

I went through the install manual again making sure I did not miss anything and from what I can see everything is set.

Jim

Subject: Re: Iptables question
Posted by dev on Fri, 31 Mar 2006 09:05:45 GMT
View Forum Message <> Reply to Message

Probably I had to ask you what host OS do you use first?
Can you provide a remote access to this node? If yes, just send me a private message or email dev at openvz dot org.
If not, it's ok, I will suggest some other ideas to check.
