Subject: *SOLVED* OpenVZ and Bastille/iptables?
Posted by marsvin on Fri, 22 Dec 2006 23:17:58 GMT
View Forum Message <> Reply to Message

Hi guys,

I've been trying out OpenVZ for the first time this week and so far it's been really easy and fun to play with. But then I decided I needed to secure my system a bit and I installed Bastille on VE0.

VE0 itself still works great but the other VEs have become completely inaccessible to all outside connections except directly from VE0.

It makes sense that this would require some extra configuration but I have no idea where to start (other than to list venet+ in /etc/Bastille/firewall.conf). Even Google turned up nothing Can anyone here point me in the right direction?

Oh I did check the routes and sysctl.config and everything looked the same as before Bastille was installed. Also flushing all rules (and replacing them with allow all) doesn't make any difference.

-- marsvin

Subject: Re: OpenVZ and Bastille/iptables?
Posted by Vasily Tarasov on Sat, 23 Dec 2006 07:29:27 GMT
View Forum Message <> Reply to Message

Hello,

Please post here the output of `route -nv`, `iptables -nv -L`, `ifconfig -a`
commands in VE/VE0 and `cat /proc/sys/net/ipv4/ip_forwarding` output in VE0.

Also, I would say, if you're able to provide remote access to your node,
it'll be much quicker to solve the problem. If it's possible, send login
information via PM.

Thanks,
Vasily.

Subject: Re: OpenVZ and Bastille/iptables?
Posted by marsvin on Sat, 23 Dec 2006 19:22:05 GMT
View Forum Message <> Reply to Message

Ah.. Actually you already helped me find it.

IP forwarding, although set to 1 in sysctl.conf for some reason was still disabled in VE0. So the fix was as simple as

echo 1 > /proc/sys/net/ipv4/ip_forward

I'm guessing Bastille disabled it at some point (although I'm fairly sure it never asked me.)

Thanks a ton for your reply you saved me a lot of trouble

Subject: Re: OpenVZ and Bastille/iptables?
Posted by marsvin on Sat, 23 Dec 2006 22:39:00 GMT
View Forum Message <> Reply to Message

You know I think I jumped the gun a bit there

Enabling ipforwarding manually made the VPS work again but when the firewall is running it still blocks all access to the VPS. At least it seems to block the routing to the VPS, since no dropped packets turn up in the log I just get this:

ssh 10.0.0.102
ssh: connect to host 10.0.0.102 port 22: No route to host

I will list the things you asked for (with the firewall enabled):

route -nv

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
10.0.0.102	0.0.0.0	255.255.255.255	UH	0	0	0	venet0
10.0.0.103	0.0.0.0	255.255.255.255	UH	0	0	0	venet0
10.0.0.101	0.0.0.0	255.255.255.255	UH	0	0	0	venet0
10.0.0.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
0.0.0.0	10.0.0.2	0.0.0.0	UG	0	0	0	eth0

ifconfig -a

```
eth0      Link encap:Ethernet  HWaddr 00:30:48:5C:28:60
          inet addr:10.0.0.12  Bcast:255.255.255.255  Mask:255.255.255.0
          inet6 addr: fe80::230:48ff:fe5c:2860/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:14344 errors:0 dropped:0 overruns:0 frame:0
          TX packets:10559 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1415554 (1.3 MiB)  TX bytes:1706761 (1.6 MiB)
          Interrupt:177

eth1      Link encap:Ethernet  HWaddr 00:30:48:5C:28:61
          BROADCAST MULTICAST  MTU:1500  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:1000
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
          Interrupt:193

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)

sit0      Link encap:IPv6-in-IPv4
          NOARP  MTU:1480  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
          UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
          RX packets:5618 errors:0 dropped:0 overruns:0 frame:0
          TX packets:5561 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:535413 (522.8 KiB)  TX bytes:523141 (510.8 KiB)


cat /proc/sys/net/ipv4/ip_forwarding
1

iptables -nv -L
Chain INPUT (policy DROP 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source               destination
    0     0 LOG        tcp  --  !lo    *       0.0.0.0/0           127.0.0.0/8         LOG flags 0 level 4 prefix
`INPUT DROP 0'
    0     0 DROP       tcp  --  !lo    *       0.0.0.0/0           127.0.0.0/8
   47  3052 ACCEPT     0    --  *      *       0.0.0.0/0           0.0.0.0/0           state
RELATED,ESTABLISHED
    0     0 LOG        0    -f  *      *       0.0.0.0/0           0.0.0.0/0           LOG flags 0 level 4 prefix
`INPUT DROP 1 '
    0     0 DROP       0    -f  *      *       0.0.0.0/0           0.0.0.0/0
    0     0 ACCEPT     0    --  lo     *       0.0.0.0/0           0.0.0.0/0
    0     0 LOG        0    --  *      *       224.0.0.0/4         0.0.0.0/0           LOG flags 0 level 4 prefix
`INPUT DROP 2 '
    0     0 DROP       0    --  *      *       224.0.0.0/4         0.0.0.0/0
```

```
    0     0 PUB_IN     0    --  eth+   *       0.0.0.0/0           0.0.0.0/0
    0     0 PUB_IN     0    --  ppp+   *       0.0.0.0/0           0.0.0.0/0
    0     0 PUB_IN     0    --  slip+  *       0.0.0.0/0           0.0.0.0/0
    0     0 DROP       tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpts:137:139
    0     0 DROP       udp  --  *      *       0.0.0.0/0           0.0.0.0/0           udp dpts:137:139
    0     0 DROP       0    --  *      *       0.0.0.0/0           224.0.0.0/8
    0     0 LOG        0    --  *      *       0.0.0.0/0           0.0.0.0/0           LOG flags 0 level 4 prefix
`INPUT DROP 4 '
    0     0 DROP       0    --  *      *       0.0.0.0/0           0.0.0.0/0

Chain FORWARD (policy DROP 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     0    --  *      *       0.0.0.0/0           0.0.0.0/0           state
RELATED,ESTABLISHED

Chain OUTPUT (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source               destination
   37 13716 PUB_OUT    0    --  *      eth+    0.0.0.0/0           0.0.0.0/0
    0     0 PUB_OUT    0    --  *      ppp+    0.0.0.0/0           0.0.0.0/0
    0     0 PUB_OUT    0    --  *      slip+   0.0.0.0/0           0.0.0.0/0

Chain INT_IN (0 references)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     icmp --  *      *       0.0.0.0/0           0.0.0.0/0
    0     0 DROP       tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpts:137:139
    0     0 DROP       udp  --  *      *       0.0.0.0/0           0.0.0.0/0           udp dpts:137:139
    0     0 DROP       0    --  *      *       0.0.0.0/0           224.0.0.0/8
    0     0 LOG        0    --  *      *       0.0.0.0/0           0.0.0.0/0           LOG flags 0 level 4 prefix
`INT_IN DROP 6 '
    0     0 DROP       0    --  *      *       0.0.0.0/0           0.0.0.0/0

Chain INT_OUT (0 references)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     icmp --  *      *       0.0.0.0/0           0.0.0.0/0
    0     0 ACCEPT     0    --  *      *       0.0.0.0/0           0.0.0.0/0

Chain PAROLE (9 references)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     0    --  *      *       0.0.0.0/0           0.0.0.0/0

Chain PUB_IN (3 references)
 pkts bytes target     prot opt in     out     source               destination
    0     0 ACCEPT     icmp --  *      *       0.0.0.0/0           0.0.0.0/0           icmp type 3
    0     0 ACCEPT     icmp --  *      *       0.0.0.0/0           0.0.0.0/0           icmp type 0
    0     0 ACCEPT     icmp --  *      *       0.0.0.0/0           0.0.0.0/0           icmp type 11
    0     0 PAROLE     tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:22
    0     0 PAROLE     tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:80
    0     0 PAROLE     tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:443
```

```
    0     0 PAROLE    tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:25
    0     0 PAROLE    tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:110
    0     0 PAROLE    tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:143
    0     0 PAROLE    tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:465
    0     0 PAROLE    tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:993
    0     0 PAROLE    tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpt:995
    0     0 DROP      tcp  --  *      *       0.0.0.0/0           0.0.0.0/0           tcp dpts:137:139
    0     0 DROP      udp  --  *      *       0.0.0.0/0           0.0.0.0/0           udp dpts:137:139
    0     0 DROP      0    --  *      *       0.0.0.0/0           224.0.0.0/8
    0     0 LOG       icmp --  *      *       0.0.0.0/0           0.0.0.0/0           LOG flags 0 level 4 prefix
`PUB_IN DROP 3'
    0     0 DROP      icmp --  *      *       0.0.0.0/0           0.0.0.0/0
    0     0 LOG       0    --  *      *       0.0.0.0/0           0.0.0.0/0           LOG flags 0 level 4 prefix
`PUB_IN DROP 5 '
    0     0 DROP      0    --  *      *       0.0.0.0/0           0.0.0.0/0

Chain PUB_OUT (3 references)
 pkts bytes target     prot opt in     out     source               destination
   35 12668 ACCEPT     0    --  *      *       0.0.0.0/0           0.0.0.0/0
```

Subject: Re: OpenVZ and Bastille/iptables?
Posted by dev on Sun, 24 Dec 2006 10:35:17 GMT

View Forum Message <> Reply to Message

I guess this is because the firewall tries to setup INPUT/OUTPUT chains, while for VPSs FORWARDING chain is working.
Maybe it has some configuration for forwarded traffic?
Another possible solution is to setup firewall inside each VE separately.

Subject: Re: OpenVZ and Bastille/iptables?
Posted by marsvin on Sun, 24 Dec 2006 15:18:08 GMT

View Forum Message <> Reply to Message

I just tried messing around with it some more and now it works! It turned out simply adding forward/accept rules to each of the VPS was necessary. Thanks a lot for keeping me on track guys.

For anyone else reading this and looking for the same thing, I created this file: /etc/Bastille/firewall.d/post-rule-setup.sh (in VE0) and added these lines:

```
iptables -A FORWARD -p tcp -d 10.0.0.101 --dport 22 --syn -j ACCEPT
iptables -A FORWARD -s 10.0.0.101 -j ACCEPT
```

This is accepting all outbound connections from my VPS and inbound for ssh. Very nice

-- marsvin
