Subject: iptables modules in VE
Posted by crea on Fri, 22 Aug 2008 12:35:00 GMT
View Forum Message <> Reply to Message

My setup:
Debian (etch-lenny mixed kind) , kernel 2.6.18 with patch-ovz028stab053.14-combined
also addon modules installed (xtables-addons-1.5.5)

On VE:
vps:/# shorewall show capabilities | grep Not
  Ipset Match: Not available
  CONNMARK Target: Not available
  Connmark Match: Not available
  Raw Table: Not available
  IPP2P Match: Not available

On HW Node:
node:~# shorewall show capabilities | grep Not
  CONNMARK Target: Not available
  Connmark Match: Not available
  IPP2P Match: Not available

All stuff about ipset is listed in /etc/vz/vz.conf (IPTABLES=..).
Does it work in VE at all ?

Second question: what do I need IPTABLES="..." for in /etc/vz/vz.conf when vzctl manual clearly says "by default all iptables modules that are loaded in the host system are accessible inside a VE". Would it be enough to put modules I need in VE in node's /etc/modules and hope it works ?

Subject: Re: iptables modules in VE
Posted by khorenko on Fri, 22 Aug 2008 12:52:28 GMT
View Forum Message <> Reply to Message

Hello.
Quote:
Second question: what do I need IPTABLES="..." for in /etc/vz/vz.conf when vzctl manual clearly says "by default all iptables modules that are loaded in the host system are accessible inside a VE". Would it be enough to put modules I need in VE in node's /etc/modules and hope it works ?

That means, if you do not have IPTABLES variable in CT.conf, all iptables modules loaded before the Container start - will be available inside a Container.

Quote:Ipset Match: Not available
Raw Table: Not available
First of all - can you please check that Container restart does not help? i mean - corresponding modules could be loaded after the Container already started, then the modules will be available on the Host System, but not inside a Container.

Second - ok, probably that modules are simply not virtualized yet. Do you really need their functionality of just was curious in general?

Thank you.

--
Konstantin

Subject: Re: iptables modules in VE
Posted by crea on Fri, 22 Aug 2008 13:21:59 GMT
View Forum Message <> Reply to Message

So am I right in understanding I can just use IPTABLES="" in /etc/vz/vz.conf since my iptables modules will be available at node startup (modules in /etc/modules will be loaded before any VE is started) ?

Yes I did 'vzctl restart 101' several times (101 is VE I experiment with).
Now I put IPTABLES="" in /etc/vz/vz.conf, did 'vzctl restart 101'
and result is:

```
vps:/# shorewall show capabilities | grep Not
  NAT: Not available
  Ipset Match: Not available
  CONNMARK Target: Not available
  Connmark Match: Not available
  Raw Table: Not available
  IPP2P Match: Not available
```

Should I just comment IPTABLES line out completely ? Why NAT became unavailable ?
And it didn't help to fix Ipset as you see anyway.

Subject: Re: iptables modules in VE
Posted by khorenko on Fri, 22 Aug 2008 13:46:42 GMT
View Forum Message <> Reply to Message

crea wrote on Fri, 22 August 2008 17:21So am I right in understanding I can just use IPTABLES="" in /etc/vz/vz.conf since my iptables modules will be available at node startup (modules in /etc/modules will be loaded before any VE is started) ?
Well, not exactly.
You need leave IPTABLES as is in /etc/vz/vz.conf and comment out completely from CT.conf (for example /etc/vz/conf/101.conf).
Then all iptables modules loaded before CT start will be available inside a Container.

Quote:Should I just comment IPTABLES line out completely ? Why NAT became unavailable ?
Because you removed IPTABLES variable from global config /etc/vz/vz.conf.

Quote:And it didn't help to fix Ipset as you see anyway.
Well, and could you please tell me if you really need the functionality of Ipset module or just curious why it's not available?

--
Konstantin

Subject: Re: iptables modules in VE
Posted by crea on Fri, 22 Aug 2008 13:53:05 GMT
View Forum Message <> Reply to Message

Ofcourse I can live without ipset, but it would be sad. I read that ipset makes matching very fast so you can have long matching rules, blacklists for example, without significant performance loss. If I can choose, I just use solution that wouldn't make me any headache in future, so I would use ipset and don't care that long iptables rules slow my server.
I can make my firewall reside completely in HW node but that is not proper design (everything about VE should be in VE)

Subject: Re: iptables modules in VE
Posted by khorenko on Fri, 22 Aug 2008 14:35:06 GMT
View Forum Message <> Reply to Message

Ok, we'll check it.
http://bugzilla.openvz.org/show_bug.cgi?id=977

Subject: Re: iptables modules in VE
Posted by maratrus on Thu, 28 Aug 2008 08:20:17 GMT
View Forum Message <> Reply to Message

Hi,

try the following way:

- make sure that ipset utility is installed inside VE
- make sure that iptables utility knows about "set" module for example:

#iptables -m set --help

- if the previous points are done try to do the following:

#ipset -N test iphash
#ipset -X iphash

if the first command fails, you have to give your VE net_admin capability:

#vzctl stop VE_ID
#vzctl set VE_ID --capability net_admin:on --save
#vzctl start VE_ID

where VE_ID - is an ID of your VE.
after that the previous ipset commands inside VE shoud work.
- make sure that you are able to use ipset module inside VE:

ipset -N mytest iphash
iptables -A FORWARD -m set --set mytest src -j ACCEPT
iptables -D FORWARD -m set --set mytest src -j ACCEPT
ipset -X mytest

if this test is success (it issues without errors) the command

shorewall show capabilities

should show Ipset Match: Available inside VE.

P.S. But keep in mind that the group of ip_set modules are not virtualized, so all of yours VEs and HN use the same resources and this is the violation of encapsulation.
Also be careful with permitting various capabilities to your VE.

P.P.S. I'm afraid that these modules won't be virtualized right now, because ipset modules are not included in mainstream kernel and goes like the extensions.