Subject: nat , ip tables, eth0 issue
Posted by marcin4 on Thu, 27 Feb 2014 21:07:44 GMT
View Forum Message <> Reply to Message

I have openvz server with bunch of containers, some on public, some on private IPs.
The issue I am having is; after week or so the container with private IP is loosing routing, but not completely.
Some packages are getting to private IP container and some do not.
monitoring of incoming packages that are not getting to cointainer showing hardware node responding "unreachable"
this is my iptables setup:
#internet access to containers
iptables -t nat -A POSTROUTING -s 10.0.1.0/24 -o eth0 -j SNAT --to A.B.C.D
iptables -A INPUT -s 10.0.1.0/24 -j ACCEPT

iptables -t nat -A PREROUTING -p tcp -d A.B.C.D --dport 8082 -i eth0 -j DNAT --to-destination 10.0.1.2:80 #web
iptables -t nat -A PREROUTING -p udp -d A.B.C.D --dport 5064 -i eth0 -j DNAT --to-destination 10.0.1.2:5064 #sip
iptables -t nat -A PREROUTING -p udp -d A.B.C.D --dport 10100:10199 -i eth0 -j DNAT --to-destination 10.0.1.2 #rtp

where A.B.C.D is a hardware node external IP address on eth0

ip route flush cache does nothing to help, nor does the restart of the effected containers.
flushing iptables and reapplying rules does nothing.
The issue must be on hardware node. MTU size on eth0 perhaps?

The only thing that helps so far is restarting network service on hardware node, short of rebooting it.
dmesg nor logs shows no errors

The public IP containers are working fine

Subject: Re: nat , ip tables, eth0 issue
Posted by Paparaciz on Fri, 28 Feb 2014 12:39:44 GMT
View Forum Message <> Reply to Message

please report your os, kernel and vz tool versions. both (HN and CT)

Subject: Re: nat , ip tables, eth0 issue
Posted by marcin4 on Fri, 28 Feb 2014 15:40:00 GMT
View Forum Message <> Reply to Message

HN CentOS release 5.10 (Final)
CT CentOS release 5.10 (Final)
2.6.18-348.16.1.el5.028stab108.1PAE
vzctl version 4.6.1

Subject: Re: nat , ip tables, eth0 issue
Posted by Paparaciz on Fri, 28 Feb 2014 18:58:20 GMT
View Forum Message <> Reply to Message

Hm, strange...

maybe some changes were made?
is same behavior with older kernel?
or can you try newer kernel?

personaly I have few HN's with centos5, kernel varies from yours (some are older some are 109.2PAE) and don't see such behavior. only my setup differs- in this HN's I only use NAT, no public ip addresses for CT's

Subject: Re: nat iptables SNAT problem
Posted by marcin4 on Sun, 02 Mar 2014 21:44:41 GMT
View Forum Message <> Reply to Message

It been always happening, even with older kernel.
I did more checking and tracing of the packages. this is what I discovered, but first let me describe the topography of my network.

block of public ips and number of servers on it.
openvz server has one of public ips,
some containers on openvz HN also have public ip addresses and
a bunch of containers in the same openvz HN have private ips netted by the same HN.
Since masquerade do not work with openvz I am using -j SNAT --to-source $WIP where $WIP is public IP address of HN.

The issue is, again not always, when the private IP containers try to access other servers on same block of public IPs.
the trace shows that these other servers see the packages coming from private IP address not from public IP of HN.
The accessing any other resource from outside my network works just fine and all the time.
CT with private IP send the UDP package to server A on public IP.
Server A sees the private IP as a source of the package.

The conclusion is that the iptables SNAT is not rewriting the source IP :

iptables -t nat -A POSTROUTING -s 10.0.1.0/24 -o $IF1 -j SNAT --to-source $WIP
where $IF1 is a eth device with public IP of HD and $WIP is the public IP of HN

If I stop the private IP CT for 30 sec or so and then restart it, everything works as it should for while (days)

I probably made this too complicated, but please try to follow my logic.
Thank you in advance