
Subject: SNAT/MASQUERADE

Posted by [sHaggY_caT](#) on Sun, 20 Dec 2009 13:23:26 GMT

[View Forum Message](#) <> [Reply to Message](#)

SNAT/MASQUERADE

http://wiki.openvz.org/Using_NAT_for_VE_with_private_IPs

```
[root@ovz-test2 ~]# vzlist | grep 407
 407      14 running 10.0.5.47    test-dns.local
[root@ovz-test2 ~]#
```

ICMP from HN to VE:

```
[root@ovz-test2 ~]# ping -c 1 10.0.5.47
PING 10.0.5.47 (10.0.5.47) 56(84) bytes of data.
64 bytes from 10.0.5.47: icmp_seq=1 ttl=64 time=0.258 ms
```

```
--- 10.0.5.47 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.258/0.258/0.258/0.000 ms
[root@ovz-test2 ~]#
```

ICMP from VE to HN:

```
[root@test-dns ~]# ping -c 1 ovz-test2
PING ovz-test2.local (10.0.5.128) 56(84) bytes of data.
64 bytes from ovz-test2.local (10.0.5.128): icmp_seq=1 ttl=64 time=0.064 ms
```

```
--- ovz-test2.local ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.064/0.064/0.064/0.000 ms
[root@test-dns ~]#
```

And, icmp from VE to another host in LAN:

```
[root@test-dns ~]# ping -c 1 puppet
PING puppet.local (10.0.5.16) 56(84) bytes of data.
64 bytes from puppet.loc (10.0.5.16): icmp_seq=1 ttl=63 time=1.78 ms
```

```
--- puppet.local ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.780/1.780/1.780/0.000 ms
[root@test-dns ~]#
```

But, NAT to another networks, for example for internet, doesn't work:

```
[root@test-dns ~]# ping -c 1 google.com
PING google.com (74.125.77.147) 56(84) bytes of data.
From ovz-test2.local (10.0.5.128) icmp_seq=1 Destination Net Unreachable
```

```
--- google.com ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms
```

```
[root@test-dns ~]#
```

```
[root@test-dns ~]# wget google.com
--2009-12-20 13:11:19-- http://google.com/
Resolving google.com... 74.125.77.104, 74.125.77.99, 74.125.77.147
Connecting to google.com[74.125.77.104]:80... failed: Network is unreachable.
```

```
Connecting to google.com[74.125.77.99]:80... failed: Network is unreachable.
Connecting to google.com[74.125.77.147]:80... failed: Network is unreachable.
[root@test-dns ~]#
```

Configuration of HN:

```
[root@ovz-test2 ~]# cat /etc/redhat-release
CentOS release 5.3 (Final)
[root@ovz-test2 ~]#
```

```
[root@ovz-test2 ~]# uname -a
Linux ovz-test2.local 2.6.18-128.2.1.el5.028stab064.4 #1 SMP Wed Jul
22 00:11:00 MSD 2009 i686 i686 i386 GNU/Linux
[root@ovz-test2 ~]#
```

```
[root@ovz-test2 ~]# ifconfig
eth0    Link encap:Ethernet  HWaddr 54:52:00:3D:CB:40
        inet addr:10.0.5.128  Bcast:10.0.5.255  Mask:255.255.255.0
        inet6 addr: fe80::5652:ff:fe3d:cb40/64  Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:112743 errors:0 dropped:0 overruns:0 frame:0
        TX packets:119926 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:21101421 (20.1 MiB)  TX bytes:23473181 (22.3 MiB)
        Interrupt:11 Base address:0x4000
```

```
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: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:878 (878.0 b)  TX bytes:878 (878.0 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:267 errors:0 dropped:0 overruns:0 frame:0
        TX packets:368 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:28529 (27.8 KiB)  TX bytes:29631 (28.9 KiB)
```

```
[root@ovz-test2 ~]#
```

```
[root@ovz-test2 ~]# rpm -qa | grep vz
vzctl-lib-3.0.23-1
```

```
vzrpm43-python-4.3.3-7_nonptl.6
vzrpm44-4.4.1-22.5
vzrpm43-4.3.3-7_nonptl.6
vzquota-3.0.12-1
vzpkg-2.7.0-18
ovzkernel-2.6.18-128.2.1.el5.028stab064.4
vzrpm44-python-4.4.1-22.5
vzctl-3.0.23-1
vzdump-1.1-2
vzyum-2.4.0-11
ha-ovz-tools-1.2-1
[root@ovz-test2 ~]#
```

```
[root@ovz-test2 ~]# sysctl -p
net.ipv4.conf.default.forwarding = 1
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.ip_forward = 1
net.ipv4.conf.all.rp_filter = 1
kernel.sysrq = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
[root@ovz-test2 ~]#
```

```
[root@ovz-test2 ~]# cat /etc/sysconfig/vz
## Global parameters
VIRTUOZZO=yes
LOCKDIR=/vz/lock
DUMPDIR=/vz/dump
VE0CPUUNITS=1000
```

```
## Logging parameters
LOGGING=yes
LOGFILE=/var/log/vzctl.log
LOG_LEVEL=0
VERBOSE=0
```

```
## Disk quota parameters
DISK_QUOTA=yes
VZFASTBOOT=no
```

```
# Disable module loading. If set, vz initscript do not load any modules.
#MODULES_DISABLED=yes
```

```
# The name of the device whose IP address will be used as source IP for CT.
# By default automatically assigned.
VE_ROUTE_SRC_DEV="eth0"
```

```
# Controls which interfaces to send ARP requests and modify APR tables on.
```

```
NEIGHBOUR_DEVS=detect
```

```
## Template parameters  
TEMPLATE=/vz/template
```

```
## Defaults for containers  
VE_ROOT=/vz/root/$VEID  
VE_PRIVATE=/vz/private/$VEID  
CONFIGFILE="vps.basic"  
DEF_OSTEMPLATE="fedora-core-4"
```

```
## Load vzwdog module  
VZWDOG="no"
```

```
## IPv4 iptables kernel modules  
IPTABLES="iptables iptable_filter iptable_mangle ipt_limit ipt_multiport  
ipt_tos ipt_TOS ipt_REJECT ipt_TCPMSS ipt_tcpmss ipt_ttl ipt_LOG  
ipt_length ip_conntrack ip_conntrack_ftp ip_conntrack_irc  
ipt_conntrack ipt_state ipt_helper iptable_nat ip_nat_ftp ip_nat_irc  
ipt_REDIRECT"
```

```
## Enable IPv6  
IPV6="no"
```

```
## IPv6 ip6tables kernel modules  
IP6TABLES="ip6_tables ip6table_filter ip6table_mangle ip6t_REJECT"
```

```
[root@ovz-test2 ~]# cat /proc/sys/net/ipv4/conf/eth0/forwarding  
1  
[root@ovz-test2 ~]# cat /proc/sys/net/ipv4/conf/venet0/forwarding  
1  
[root@ovz-test2 ~]#
```

```
[root@ovz-test2 ~]# iptables-save  
# Generated by iptables-save v1.3.5 on Sun Dec 20 13:17:25 2009  
*raw  
:PREROUTING ACCEPT [9708:1526221]  
:OUTPUT ACCEPT [9198:1571058]  
COMMIT  
# Completed on Sun Dec 20 13:17:25 2009  
# Generated by iptables-save v1.3.5 on Sun Dec 20 13:17:25 2009  
*nat  
:PREROUTING ACCEPT [73:4765]  
:POSTROUTING ACCEPT [0:0]
```

```

:OUTPUT ACCEPT [945:55800]
-A POSTROUTING -m mark --mark 0x9 -j MASQUERADE
-A POSTROUTING -o eth0 -j MASQUERADE
-A POSTROUTING -o venet0 -j MASQUERADE
-A POSTROUTING -m mark --mark 0x9 -j MASQUERADE
COMMIT
# Completed on Sun Dec 20 13:17:25 2009
# Generated by iptables-save v1.3.5 on Sun Dec 20 13:17:25 2009
*mangle
:PREROUTING ACCEPT [11775:1810121]
:INPUT ACCEPT [11090:1747639]
:FORWARD ACCEPT [668:61270]
:OUTPUT ACCEPT [11071:1902912]
:POSTROUTING ACCEPT [11739:1964182]
-A PREROUTING -i eth0 -j MARK --set-mark 0x9
-A PREROUTING -i venet0 -j MARK --set-mark 0x9
-A PREROUTING -i eth0 -j MARK --set-mark 0x9
-A PREROUTING -i venet0 -j MARK --set-mark 0x9
COMMIT
# Completed on Sun Dec 20 13:17:25 2009
# Generated by iptables-save v1.3.5 on Sun Dec 20 13:17:25 2009
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [11071:1902912]
:RH-Firewall-1-INPUT - [0:0]
-A INPUT -j RH-Firewall-1-INPUT
-A FORWARD -j RH-Firewall-1-INPUT
-A RH-Firewall-1-INPUT -i lo -j ACCEPT
-A RH-Firewall-1-INPUT -i eth0 -j ACCEPT
-A RH-Firewall-1-INPUT -i venet0 -j ACCEPT
-A RH-Firewall-1-INPUT -p icmp -m icmp --icmp-type any -j ACCEPT
-A RH-Firewall-1-INPUT -p esp -j ACCEPT
-A RH-Firewall-1-INPUT -p ah -j ACCEPT
-A RH-Firewall-1-INPUT -d 224.0.0.251 -p udp -m udp --dport 5353 -j ACCEPT
-A RH-Firewall-1-INPUT -p udp -m udp --dport 631 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m tcp --dport 631 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A RH-Firewall-1-INPUT -j REJECT --reject-with icmp-host-prohibited
COMMIT
# Completed on Sun Dec 20 13:17:25 2009

```

```

[root@ovz-test2 ~]# 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  all  -- anywhere      anywhere
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
REJECT  all  -- anywhere      anywhere
reject-with icmp-host-prohibited
[root@ovz-test2 ~]#
```

=====

*nat

```
-A POSTROUTING -m mark --mark 0x9 -j MASQUERADE
```

*mangle

```
POSTROUTING ACCEPT [15696772:11135009050]
```

```
-A PREROUTING -i br0 -j MARK --set-mark 0x9
```

```
-A PREROUTING -i wlan0 -j MARK --set-mark 0x9
```

```
-A PREROUTING -i venet0 -j MARK --set-mark 0x9
```

```
-A FORWARD -o eth0 -j ACCEPT
```

```
-A RH-Firewall-1-INPUT -i br0 -j ACCEPT
```

```
-A RH-Firewall-1-INPUT -i ath0 -j ACCEPT
```

```
-A RH-Firewall-1-INPUT -i venet0 -j ACCEPT
```

Subject: Re: SNAT/MASQUERADE
Posted by [sHaggY_caT](#) on Sun, 20 Dec 2009 19:33:03 GMT
[View Forum Message](#) <> [Reply to Message](#)
