
Subject: *SOLVED* NAT doesn't seem to work
Posted by [duswil](#) on Tue, 08 Aug 2006 00:55:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

I have followed the instructions at [Using_NAT_for_VE_with_private_IPs](#).

My kernel does *not* have loadable modules enabled. Even with that enabled it doesn't seem to work.

I'm trying to have a set of VEs that have only an internal IP address (172.16.*.*) that can access remote services (apt-get, ping, etc). I don't want these VEs to have public IP addresses. Eventually they will be accessed through a proxy server in a VE that does have a public IP. It is important that these non-public VEs have access to Internet services.

The hn has Red Hat Enterprise Linux ES 4 with a custom compiled kernel. Attached is my .config file.

```
[root@hn ~]# uname -a
Linux hn.fakedomainname.com 2.6.8-022stab078-up #15 Tue Aug 8 00:13:33 UTC 2006 i686
athlon i386 GNU/Linux
```

```
[root@hn ~]# cat /etc/modules.conf
options ip_contrack ip_contrack_enable_ve0=1
```

The 27.16.122.72 IP is fake, I used my real HN IP address in its place:
iptables -t nat -A POSTROUTING -s 172.16.0.0/16 -o eth0 -j SNAT --to 27.16.122.72

"testvps" has a private IP address:

```
testvps:~# ifconfig
```

```
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:1397 errors:0 dropped:0 overruns:0 frame:0
        TX packets:1397 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:111809 (109.1 KiB)  TX bytes:111809 (109.1 KiB)
```

```
venet0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
        inet addr:127.0.0.1  P-t-P:127.0.0.1  Bcast:0.0.0.0  Mask:255.255.255.255
        UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
        RX packets:2 errors:0 dropped:0 overruns:0 frame:0
        TX packets:26 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:480 (480.0 b)  TX bytes:2128 (2.0 KiB)
```

```
venet0:0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
        inet addr:172.16.4.1  P-t-P:172.16.4.1  Bcast:0.0.0.0  Mask:255.255.255.255
        UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
```

```
Problem with pinging google.com from "testvps":
testvps:~# ping -w 10 google.com
PING google.com (64.233.167.99) 56(84) bytes of data.
```

```
--- google.com ping statistics ---
11 packets transmitted, 0 received, 100% packet loss, time 9998ms
```

```
"dns" has a public IP (replaced with a fake one):
```

```
dns:~# ifconfig
```

```
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:1462 errors:0 dropped:0 overruns:0 frame:0
      TX packets:1462 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:116801 (114.0 KiB)  TX bytes:116801 (114.0 KiB)
```

```
venet0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      inet addr:127.0.0.1  P-t-P:127.0.0.1  Bcast:0.0.0.0  Mask:255.255.255.255
      UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
      RX packets:57 errors:0 dropped:0 overruns:0 frame:0
      TX packets:56 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:12536 (12.2 KiB)  TX bytes:7094 (6.9 KiB)
```

```
venet0:0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      inet addr:27.16.122.128  P-t-P:27.16.122.128  Bcast:0.0.0.0  Mask:255.255.255.255
      UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
```

```
venet0:1  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      inet addr:172.16.1.1  P-t-P:172.16.1.1  Bcast:0.0.0.0  Mask:255.255.255.255
      UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
```

```
No problem with pinging google.com from "dns":
```

```
dns:~# ping -w 2 google.com
```

```
PING google.com (64.233.187.99) 56(84) bytes of data.
```

```
64 bytes from 64.233.187.99: icmp_seq=1 ttl=239 time=37.9 ms
```

```
64 bytes from 64.233.187.99: icmp_seq=2 ttl=239 time=38.1 ms
```

```
--- google.com ping statistics ---
```

```
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
```

```
rtt min/avg/max/mdev = 37.993/38.049/38.106/0.203 ms
```

Let me know what else might help track this problem down. I just can't seem figure it out on my own.

Thanks!

Dusty

File Attachments

1) [dot.config](#), downloaded 633 times

Subject: Re: NAT doesn't seem to work

Posted by [duswil](#) on Tue, 08 Aug 2006 01:33:43 GMT

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

```
[root@hn net]# cat /proc/sys/net/ipv4/ip_forward
1
```

```
[root@hn net]# cat /etc/sysconfig/vz-scripts/4001.conf
# Configuration file generated by vzsplit for 20 VPS
# on HN with total amount of physical mem 2023 Mb
# low memory 2023 Mb, swap size 1027 Mb, Max treads 8000
# Resource commit level 0:
# Free resource distribution. Any parameters may be increased
# Primary parameters
NUMPROC="516:516"
AVNUMPROC="258:258"
NUMTCPSOCK="516:516"
NUMOTHERSOCK="516:516"
VMGUARPAGES="28052:2147483647"
```

```
# Secondary parameters
KMEMSIZE="21216542:23338196"
TCPSNDBUF="4958645:7072181"
TCPRCVBUF="4958645:7072181"
OTHERSOCKBUF="2479322:4592858"
DGRAMRCVBUF="2479322:2479322"
OOMGUARPAGES="28052:2147483647"
PRIVVMPAGES="168312:185143"
```

```
# Auxiliary parameters
LOCKEDPAGES="1035:1035"
SHMPAGES="16831:16831"
PHYSPAGES="0:2147483647"
NUMFILE="8256:8256"
NUMFLOCK="825:907"
NUMPTY="51:51"
NUMSIGINFO="1024:1024"
DCACHESIZE="4616947:4755456"
NUMIPTENT="100:100"
DISKSPACE="724049:796454"
DISKINODES="396156:435772"
CPUUNITS="4680"
```

```
IP_ADDRESS="172.16.4.1"
HOSTNAME="testvps"
VE_ROOT="/vz/root/$VEID"
VE_PRIVATE="/vz/private/$VEID"
OSTEMPLATE="debian-3.1-i386-base"
ORIGIN_SAMPLE="default"
NAMESERVER="172.16.1.1"
SEARCHDOMAIN="local"
ONBOOT="yes"
```

```
[root@hn net]# cat /etc/vz/vz.conf
```

```
## Global parameters
```

```
VIRTUOZZO=yes
```

```
LOCKDIR=/vz/lock
```

```
DUMPDIR=/vz/dump
```

```
VE0CPUUNITS=1000
```

```
## Logging parameters
```

```
LOGGING=yes
```

```
LOGFILE=/var/log/vzctl.log
```

```
LOG_LEVEL=0
```

```
## Disk quota parameters
```

```
DISK_QUOTA=yes
```

```
VZFASTBOOT=no
```

```
# The name of the device whose ip address will be used as source ip for VE.
```

```
# By default automatically assigned.
```

```
#VE_ROUTE_SRC_DEV="eth0"
```

```
## Template parameters
```

```
TEMPLATE=/vz/template
```

```
## Defaults for VEs
```

```
VE_ROOT=/vz/root/$VEID
```

```
VE_PRIVATE=/vz/private/$VEID
```

```
CONFIGFILE="default"
```

```
DEF_OSTEMPLATE="debian-3.1-i386-base"
```

```
## Load vzwdog module
```

```
VZWDOG="no"
```

```
IPTABLES="ipt_REJECT ipt_tos ipt_limit ipt_multiport iptable_filter iptable_mangle ipt_TCPMSS  
ipt_tcpmss ipt_ttl ipt_length iptable_nat"
```

```
[root@hn net]# cat /etc/vz/conf/ve-default.conf-sample
```

```
# Configuration file generated by vzsplite for 20 VPS
```

```
# on HN with total amount of physical mem 2023 Mb
```

```
# low memory 2023 Mb, swap size 1027 Mb, Max treads 8000
# Resource commit level 0:
# Free resource distribution. Any parameters may be increased
# Primary parameters
NUMPROC="516:516"
AVNUMPROC="258:258"
NUMTCPSOCK="516:516"
NUMOTHERSOCK="516:516"
VMGUARPAGES="28052:2147483647"
```

```
# Secondary parameters
KMEMSIZE="21216542:23338196"
TCPSNDBUF="4958645:7072181"
TCPRCVBUF="4958645:7072181"
OTHERSOCKBUF="2479322:4592858"
DGRAMRCVBUF="2479322:2479322"
OOMGUARPAGES="28052:2147483647"
PRIVVMPAGES="168312:185143"
```

```
# Auxiliary parameters
LOCKEDPAGES="1035:1035"
SHMPAGES="16831:16831"
PHYSPAGES="0:2147483647"
NUMFILE="8256:8256"
NUMFLOCK="825:907"
NUMPTY="51:51"
NUMSIGINFO="1024:1024"
DCACHESIZE="4616947:4755456"
NUMIPTENT="100:100"
DISKSPACE="724049:796454"
DISKINODES="396156:435772"
CPUUNITS="4680"
```

Subject: Re: NAT doesn't seem to work
Posted by [duswil](#) on Tue, 08 Aug 2006 01:37:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
[root@hn net]# ip a l
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
   link/void
2: eth0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
   link/ether 00:0e:0c:83:4a:90 brd ff:ff:ff:ff:ff:ff
   inet 27.16.122.72/26 brd 27.16.122.108 scope global eth0
4: lo: <LOOPBACK,UP> mtu 16436 qdisc noqueue
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 brd 127.255.255.255 scope host lo
6: dummy0: <BROADCAST,NOARP> mtu 1500 qdisc noop
```

```
link/ether 1a:60:19:07:1a:f0 brd ff:ff:ff:ff:ff
8: tunl0: <NOARP> mtu 1480 qdisc noop
link/ipip 0.0.0.0 brd 0.0.0.0
```

```
[root@hn net]# ip r l
172.16.2.1 dev venet0 scope link src 27.16.122.72
172.16.4.1 dev venet0 scope link src 27.16.122.72
172.16.2.2 dev venet0 scope link src 27.16.122.72
27.16.122.129 dev venet0 scope link src 27.16.122.72
27.16.122.128 dev venet0 scope link src 27.16.122.72
27.16.122.131 dev venet0 scope link src 27.16.122.72
27.16.122.130 dev venet0 scope link src 27.16.122.72
172.16.1.1 dev venet0 scope link src 27.16.122.72
172.16.3.1 dev venet0 scope link src 27.16.122.72
172.16.3.3 dev venet0 scope link src 27.16.122.72
27.16.122.134 dev venet0 scope link src 27.16.122.72
27.16.122.0/26 dev eth0 proto kernel scope link src 27.16.122.72
169.254.0.0/16 dev eth0 scope link
default via 27.16.122.1 dev eth0
```

Subject: Re: NAT doesn't seem to work
Posted by [duswil](#) on Tue, 08 Aug 2006 02:09:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
[root@hn net]# cat /proc/sys/net/ipv4/conf/eth0/proxy_arp
0
```

```
[root@hn net]# cat /proc/sys/net/ipv4/conf/venet0/proxy_arp
0
```

```
[root@hn net]# cat /proc/sys/net/ipv4/conf/all/proxy_arp
0
```

```
[root@hn net]# iptables -t nat -vnL POSTROUTING
Chain POSTROUTING (policy ACCEPT 0 packets, 0 bytes)
pkts bytes target prot opt in out source destination
0 0 SNAT all -- * eth0 172.16.0.0/16 0.0.0.0/0 to:27.16.122.72
```

```
[root@hn net]# route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
172.16.2.1 * 255.255.255.255 UH 0 0 0 venet0
172.16.4.1 * 255.255.255.255 UH 0 0 0 venet0
172.16.2.2 * 255.255.255.255 UH 0 0 0 venet0
27.16.122.129 * 255.255.255.255 UH 0 0 0 venet0
27.16.122.128 * 255.255.255.255 UH 0 0 0 venet0
27.16.122.131 * 255.255.255.255 UH 0 0 0 venet0
```

```

27.16.122.130 *      255.255.255.255 UH  0  0  0 venet0
172.16.1.1 *        255.255.255.255 UH  0  0  0 venet0
172.16.3.1 *        255.255.255.255 UH  0  0  0 venet0
172.16.3.3 *        255.255.255.255 UH  0  0  0 venet0
27.16.122.134 *     255.255.255.255 UH  0  0  0 venet0
27.16.122.0 *       255.255.255.192 U   0  0  0 eth0
169.254.0.0 *       255.255.0.0   U   0  0  0 eth0
default 27.16.122.1  0.0.0.0   UG  0  0  0 eth0

```

```

[root@hn ~]# cd /proc/sys/net/ipv4/conf
[root@hn conf]# for i in */*; do echo $i; cat $i; done

```

```

all/accept_redirects
0
all/accept_source_route
0
all/arp_announce
0
all/arp_filter
0
all/arp_ignore
0
all/bootp_relay
0
all/disable_policy
0
all/disable_xfrm
0
all/force_igmp_version
0
all/forwarding
1
all/log_martians
0
all/mc_forwarding
0
all/medium_id
0
all/proxy_arp
0
all/rp_filter
0
all/secure_redirects
1
all/send_redirects
0
all/shared_media
1
all/tag

```

0
default/accept_redirects
1
default/accept_source_route
0
default/arp_announce
0
default/arp_filter
0
default/arp_ignore
0
default/bootp_relay
0
default/disable_policy
0
default/disable_xfrm
0
default/force_igmp_version
0
default/forwarding
1
default/log_martians
0
default/mc_forwarding
0
default/medium_id
0
default/proxy_arp
0
default/rp_filter
1
default/secure_redirects
1
default/send_redirects
1
default/shared_media
1
default/tag
0
eth0/accept_redirects
1
eth0/accept_source_route
0
eth0/arp_announce
0
eth0/arp_filter
0
eth0/arp_ignore

0
eth0/bootp_relay
0
eth0/disable_policy
0
eth0/disable_xfrm
0
eth0/force_igmp_version
0
eth0/forwarding
1
eth0/log_martians
0
eth0/mc_forwarding
0
eth0/medium_id
0
eth0/proxy_arp
0
eth0/rp_filter
1
eth0/secure_redirects
1
eth0/send_redirects
1
eth0/shared_media
1
eth0/tag
0
lo/accept_redirects
1
lo/accept_source_route
0
lo/arp_announce
0
lo/arp_filter
0
lo/arp_ignore
0
lo/bootp_relay
0
lo/disable_policy
0
lo/disable_xfrm
0
lo/force_igmp_version
0
lo/forwarding

1
lo/log_martians
0
lo/mc_forwarding
0
lo/medium_id
0
lo/proxy_arp
0
lo/rp_filter
1
lo/secure_redirects
1
lo/send_redirects
1
lo/shared_media
1
lo/tag
0
venet0/accept_redirects
1
venet0/accept_source_route
0
venet0/arp_announce
0
venet0/arp_filter
0
venet0/arp_ignore
0
venet0/bootp_relay
0
venet0/disable_policy
0
venet0/disable_xfrm
0
venet0/force_igmp_version
0
venet0/forwarding
1
venet0/log_martians
0
venet0/mc_forwarding
0
venet0/medium_id
0
venet0/proxy_arp
0
venet0/rp_filter

```
1
venet0/secure_redirects
1
venet0/send_redirects
0
venet0/shared_media
1
venet0/tag
0
```

```
[root@hn ~]# ifconfig
```

```
eth0  Link encap:Ethernet HWaddr 00:0E:0C:83:4A:90
      inet addr:27.16.122.72 Bcast:27.16.122.108 Mask:255.255.255.192
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:386 errors:0 dropped:0 overruns:0 frame:0
      TX packets:387 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:41379 (40.4 KiB) TX bytes:46938 (45.8 KiB)
```

```
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:417 errors:0 dropped:0 overruns:0 frame:0
      TX packets:417 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:36545 (35.6 KiB) TX bytes:36545 (35.6 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:18 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:2029 (1.9 KiB) TX bytes:1901 (1.8 KiB)
```

Subject: Re: NAT doesn't seem to work

Posted by [Vasily Tarasov](#) on Tue, 08 Aug 2006 06:08:11 GMT

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

After compiling modules in kernel, modprobe.conf isn't processed.

So the `ip_contrack_enable_ve0` parameter doesn't reach the module.

Moreover there can be other problems with using iptables compiled in.

Please, try to use kernel with loadable modules and with Network Filters compiled as modules...

Subject: Re: NAT doesn't seem to work

Posted by [duswil](#) on Tue, 08 Aug 2006 06:26:52 GMT

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

Going as modules didn't seem to make any difference. The conntrack module shows up in the lsmod list. All the same results.

Subject: Re: NAT doesn't seem to work

Posted by [duswil](#) on Tue, 08 Aug 2006 06:51:42 GMT

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

```
testvps:~# ping -w 10 google.com
PING google.com (64.233.167.99) 56(84) bytes of data.
```

```
--- google.com ping statistics ---
10 packets transmitted, 0 received, 100% packet loss, time 9087ms
```

```
[root@hn ~]# lsmod
Module                Size Used by
iptables_nat          26972 0
ip_conntrack           35720 1 iptables_nat
ipt_length             1728 0
ipt_ttl                1856 0
ipt_tcpmss             2112 0
ipt_TCPMSS             3904 0
iptables_mangle        4544 0
ipt_multiport          2048 0
ipt_limit              2176 0
ipt_tos                1600 0
ipt_REJECT             6016 0
iptables_filter        4288 1
ip_tables              20752 11
iptables_nat,ipt_length,ipt_ttl,ipt_tcpmss,ipt_TCPMSS,iptables_mangle,ipt_multiport,ipt_limit,ipt_tos,
ipt_REJECT,iptables_filter
```

As before, I replaced my real IP with 27.16.122.72.

```
[root@hn ~]# iptables -nvL
Chain INPUT (policy ACCEPT 927 packets, 78071 bytes)
pkts bytes target prot opt in out source destination
 927 78071 PORTSEN all -- * * 0.0.0.0/0 0.0.0.0/0
```

```
Chain FORWARD (policy ACCEPT 30 packets, 4722 bytes)
pkts bytes target prot opt in out source destination
```

```
Chain OUTPUT (policy ACCEPT 839 packets, 95014 bytes)
pkts bytes target prot opt in out source destination
```

```
Chain PORTSEN (1 references)
```

```
pkts bytes target  prot opt in  out  source          destination
```

```
[root@hn ~]# iptables -t nat -vnL POSTROUTING
```

```
Chain POSTROUTING (policy ACCEPT 0 packets, 0 bytes)
```

```
pkts bytes target  prot opt in  out  source          destination
 0    0 SNAT    all  --  *   eth0  172.16.0.0/16  0.0.0.0/0      to:27.16.122.72
```

Subject: Re: NAT doesn't seem to work

Posted by [duswil](#) on Tue, 08 Aug 2006 06:53:26 GMT

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

Attaching new (module-enabled) .config file.

File Attachments

1) [dot.config.txt](#), downloaded 617 times

Subject: Re: NAT doesn't seem to work

Posted by [Vasily Tarasov](#) on Tue, 08 Aug 2006 08:30:05 GMT

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

Quote:[root@hn ~]# cat /etc/modules.conf
options ip_conntrack ip_conntrack_enable_ve0=1

Please, use modprobe.conf file. And, to be sure, reboot the node after modifying the file, thanks!

Subject: Re: NAT doesn't seem to work

Posted by [duswil](#) on Tue, 08 Aug 2006 23:58:37 GMT

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

That one worked. As long as you're using modules (which I don't want to), that does the trick. Your comment helped greatly. Thanks!

I got it to work without module support, but it required patching one line of the the OpenVZ patch.

You might want to consider making this flag (ip_conntrack_enable_ve0) an option to be set in the .config (via make menuconfig, make config, or whatever). If I knew how, I would help with that.

In the patch, I found the line that said:

```
+int ip_conntrack_enable_ve0 = 0;
```

and replaced it with:

```
+int ip_conntrack_enable_ve0 = 1;
```

which hardcodes that option into the kernel.

That part of the OpenVZ patch affects the

"linux-2.6.8/net/ipv4/netfilter/ip_contrack_standalone.c" file in the kernel source.

(removed patch, see http://bugzilla.openvz.org/show_bug.cgi?id=218 instead)

Thanks!

Subject: Re: NAT doesn't seem to work
Posted by [duswil](#) on Wed, 09 Aug 2006 00:45:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

I lied. I guess I was able to add it into the menu. It was much easier than I had expected.

My proposed patch is at http://bugzilla.openvz.org/show_bug.cgi?id=218.

Subject: Re: NAT doesn't seem to work
Posted by [Vasily Tarasov](#) on Wed, 09 Aug 2006 06:00:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

Well... if you've compiled this module in you can pass parameter to it using kernel parameters. For example, using GRUB config:

```
kernel /vmlinuz root=/dev/sda2 ip_contrack_enable_ve0=1
```

Please, if possible, try this method and if it doesn't work report about it in bugzilla (bug #218).

Thank you very much!
vass

FIXED =0 -> =1. Badly missprint.

Subject: Re: NAT doesn't seem to work
Posted by [duswil](#) on Wed, 09 Aug 2006 06:04:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

That would be "=1", right?

```
kernel /vmlinuz root=/dev/sda2 ip_contrack_enable_ve0=1
```

Also, that didn't seem obvious to me at all. Maybe it would be good to document that somewhere? *shrug*

Subject: Re: NAT doesn't seem to work
Posted by [Vasily Tarasov](#) on Wed, 09 Aug 2006 06:37:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quote:Also, that didn't seem obvious to me at all. Maybe it would be good to document that somewhere? *shrug*

Just FYI, every module compiled in can receive parameters through kernel parameters. A lot of drivers do so, for example.

Of course it must be documented.

Good luck!
