
Subject: IPv6+NAT+venet IFs

Posted by [dcbun](#) on Thu, 20 Nov 2008 15:32:30 GMT

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

EDIT: Shame on me for not skimming the forums before posting. I see a half a dozen similar problems/solutions to what I seem to be experiencing; I'm going to try a few and see if I can't get this darn thing working.

Hello all, new user of the forums and OpenVZ in general.

Using debian-stable's 2.6.18-12-fza-686 kernel and vzctl (and I assume all other utilities) version 3.0.23-1dso1~etch0.

I am trying to get IPv6 connectivity working behind a NAT'd IPv4 connection, to no avail. IPv4 connectivity works fine, regardless of whether I use a venet or veth device. My router has a Hurricane Electric provided /48 delegated to it, and distributes a /64 of addresses to the rest of my LAN via radvd stateless autoconfiguration.

Originally I attempted to establish IPv6 functionality behind my NAT using static IPv6 addresses and using venet devices the same way IPv4 is done--by adding a SNAT routing rule to my iptables rules. However, (and please correct me if I am wrong), SNAT does not work with IPv6 connections.

```
root@skunk:~# ip6tables -t nat -L
```

```
ip6tables v1.3.6: can't initialize ip6tables table `nat': Table does not exist (do you need to insmod?)
```

Perhaps ip6tables or your kernel needs to be upgraded.

Such a module seems to not exist, at least in this build of my kernel. I also cannot find anything on the matter via Google, perhaps SNAT/DNAT has not yet been implemented in IPv6?

skunk=HW, wolf=VE

```
root@skunk:~# iptables -t nat -L # mangle and main table are empty, default accept
```

```
Chain PREROUTING (policy ACCEPT)
```

```
target    prot opt source                destination
```

```
Chain POSTROUTING (policy ACCEPT)
```

```
target    prot opt source                destination
```

```
SNAT     0  -- 192.168.1.0/24        anywhere           to:192.168.1.1
```

```
Chain OUTPUT (policy ACCEPT)
```

```
target    prot opt source                destination
```

```
root@skunk:~# ifconfig eth0
```

```
eth0    Link encap:Ethernet HWaddr 00:0F:B0:02:FA:F0
```

```
inet addr:192.168.1.1 Bcast:192.168.1.255 Mask:255.255.254.0
inet6 addr: 2001:470:881f:dcb:20f:b0ff:fe02:faf0/64 Scope:Global
inet6 addr: fe80::20f:b0ff:fe02:faf0/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:796 errors:0 dropped:0 overruns:0 frame:0
TX packets:543 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:86650 (84.6 KiB) TX bytes:61637 (60.1 KiB)
Interrupt:185 Base address:0x7000
```

```
root@skunk:~# vzctl create 1 --ostemplate debian-minimal
Creating container private area (debian-minimal)
Performing postcreate actions
Container private area was created
root@skunk:~# vzctl set 1 --ipadd 192.168.1.10 --hostname wolf --nameserver 192.168.0.1 --save
Saved parameters for CT 1
root@skunk:~# vzctl set 1 --ipadd 2001:0470:881f:0dcb:0000:0000:0000:1234 --save
Saved parameters for CT 1
root@skunk:~# vzctl start 1
Starting container ...
Container is mounted
Adding IP address(es): 192.168.1.10 2001:470:881f:dcb::1234
Setting CPU units: 250
Configure meminfo: 28125
Set hostname: wolf
File resolv.conf was modified
Container start in progress...
```

```
root@skunk:~# ip route list table all
192.168.1.10 dev venet0 scope link
192.168.0.0/23 dev eth0 proto kernel scope link src 192.168.1.1
default via 192.168.0.1 dev eth0
local 192.168.1.1 dev eth0 table 255 proto kernel scope host src 192.168.1.1
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
broadcast 192.168.0.0 dev eth0 table 255 proto kernel scope link src 192.168.1.1
broadcast 192.168.1.255 dev eth0 table 255 proto kernel scope link src 192.168.1.1
broadcast 127.0.0.0 dev lo table 255 proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table 255 proto kernel scope host src 127.0.0.1
local ::1 via :: dev lo proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
2001:470:881f:dcb::1234 dev venet0 metric 1024 expires 2133392sec mtu 1500 advmss 1440
hoplimit 4294967295
local 2001:470:881f:dcb:20f:b0ff:fe02:faf0 via :: dev lo proto none metric 0 mtu 16436 advmss
16376 hoplimit 4294967295
2001:470:881f:dcb::/64 dev eth0 proto kernel metric 256 expires 214714sec mtu 1500 advmss
1440 hoplimit 4294967295
local fe80::1 via :: dev lo proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
```

```
local fe80::20f:b0ff:fe02:faf0 via :: dev lo proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
fe80::/64 dev eth0 metric 256 expires 2133320sec mtu 1500 advmss 1440 hoplimit 4294967295
fe80::/64 dev venet0 metric 256 expires 2133320sec mtu 1500 advmss 1440 hoplimit 4294967295
ff02::1 via ff02::1 dev eth0 metric 0
  cache mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev eth0 metric 256 expires 2133320sec mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev venet0 metric 256 expires 2133320sec mtu 1500 advmss 1440 hoplimit 4294967295
default via fe80::200:ff:fe00:0 dev eth0 proto kernel metric 1024 expires 2sec mtu 1500 advmss 1440 hoplimit 64
unreachable default dev lo proto none metric -1 error -101 hoplimit 255
```

```
root@skunk:~# vzctl enter 1
entered into CT 1
# installed iproute via apt (thus ipv4 works)
wolf:/# ip route list table all
192.0.2.1 dev venet0 scope link
default via 192.0.2.1 dev venet0
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
local 192.168.1.10 dev venet0 table 255 proto kernel scope host src 192.168.1.10
broadcast 127.0.0.0 dev lo table 255 proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table 255 proto kernel scope host src 127.0.0.1
local ::1 via :: dev lo proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local 2001:470:881f:dcb::1234 via :: dev lo proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
fe80::/64 dev venet0 metric 256 expires 21333793sec mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev venet0 metric 256 expires 21333793sec mtu 1500 advmss 1440 hoplimit 4294967295
default dev venet0 metric 256 expires 21333793sec mtu 1500 advmss 1440 hoplimit 4294967295
```

```
wolf:/# ping6 www.kame.net
PING www.kame.net(orange.kame.net) 56 data bytes
```

```
--- www.kame.net ping statistics ---
```

```
60 packets transmitted, 0 received, 100% packet loss, time 58990ms
```

```
wolf:/# ping6 2001:200:0:8002:203:47ff:fea5:3085 # ipv6 address of www.kame.net
```

```
PING 2001:200:0:8002:203:47ff:fea5:3085(2001:200:0:8002:203:47ff:fea5:3085) 56 data bytes
```

```
--- 2001:200:0:8002:203:47ff:fea5:3085 ping statistics ---
```

```
327 packets transmitted, 0 received, 100% packet loss, time 325959ms
```

And finally, the tcpdump outputs during a ping6 to www.kame.net:

In VE1 (wolf):

```
wolf:~# tcpdump -i venet0 -e -vv -c 8 host www.kame.net
```

```
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
```

```
tcpdump: WARNING: venet0: no IPv4 address assigned
```

```
tcpdump: listening on venet0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
```

```
16:27:30.944869 Out ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 196
```

```
16:27:31.944716 Out ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 197
```

```
16:27:32.944551 Out ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 198
```

```
16:27:33.944397 Out ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 199
```

```
16:27:34.944241 Out ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 200
```

```
16:27:35.944101 Out ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 201
```

```
16:27:36.943946 Out ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 202
```

```
16:27:37.943792 Out ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 203
```

```
8 packets captured
```

```
12 packets received by filter
```

```
0 packets dropped by kernel
```

In HW node (skunk), from venet0

```
root@skunk:~# tcpdump -i venet0 -e -vv -c 8 host www.kame.net
```

```
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
```

```
tcpdump: WARNING: venet0: no IPv4 address assigned
```

```
tcpdump: listening on venet0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
```

```
11:30:03.921621 In ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 349
```

```
11:30:04.921460 In ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58), length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq 350
```

```
11:30:05.921290 In ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58),
```

```
length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq
351
11:30:06.921145 In ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58),
length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq
352
11:30:07.920985 In ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58),
length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq
353
11:30:08.920847 In ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58),
length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq
354
11:30:09.920683 In ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58),
length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq
355
11:30:10.920533 In ethertype IPv6 (0x86dd), length 120: (hlim 64, next-header: ICMPv6 (58),
length: 64) 2001:470:881f:dcb::1234 > orange.kame.net: ICMP6, echo request, length 64, seq
356
8 packets captured
8 packets received by filter
0 packets dropped by kernel
```

In HW node, from eth0 (nothing came through, so I waited about 60 seconds and ctrl+c'd it)
root@skunk:~# tcpdump -i eth0 -e -vv -c 8 host www.kame.net
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes

```
0 packets captured
4 packets received by filter
0 packets dropped by kernel
```

I had also tried this using a VE with a veth virtual ethernet interface and piggybacking off of my router's radvd advertisements. I got a bit farther progress-wise with that method but I'd like to rule out this method as an impossibility if that is the case.

Any gentle nudge in the right direction as far as fixing this will be greatly appreciated. Thanks!
