

HELP: NAT/Masquerading broken with 2.6.11 + pppoe (long)

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.networking/2005-05/0511.html>

From: Albrecht Dreß (albrecht.dress_at_arcor.de)

Date: 05/15/05

Date: Sun, 15 May 2005 19:38:42 +0200

Hi,

first, please excuse me if this is either a dumb question or the wrong forum. Any pointers would be greatly appreciated in this case, though...

I have a small home network looking as follows:

192.168.42.3

| PMac G4 | | |---DSL Modem (ppp0)
ISDN---|ipp0 eth0|---|Switch|---two computers, printer (192.168.42.x)

The machine marked as "PMac G4" is a PowerMac G4/800 "Silver", running Linux 2.6.11.4 on the Yellowdog 4.01 disto, which includes iptables v1.2.9.

I had an "old" setup with the G4 working as router via the isdn adaptor, which worked flawlessly. I now switched to ADSL, so I removed the ISDN connection and just changed to ppp0 using the kernel-based pppoe driver.

Now the machines in my "local" net can still connect the G4 via eth0 (e.g. ssh or dns; the dns runs a caching nameserver), and the G4 has full access to the internet, too. However, nat/masquerading from my local net via the G4 to the outside world fails. I stripped down my ipfilter config to a completely open one (see script below), but still had no success:

```
<snip>
#!/bin/bash
for n in INPUT OUTPUT FORWARD; do
    iptables -F $n
    iptables -P $n ACCEPT
done
for n in PREROUTING POSTROUTING OUTPUT; do
    iptables -t nat -F $n
```

comp.os.linux.networking: HELP: NAT/Masquerading broken with 2.6.11 + pppoe (long)

```
iptables -t nat -P $n ACCEPT
done
iptables -t nat -A POSTROUTING -s 192.168.42.0/24 -o ppp0 -j MASQUERADE
iptables -A FORWARD -p tcp --tcp-flags SYN,RST SYN -j TCPMSS \
  --clamp-mss-to-pmtu
echo "1" > /proc/sys/net/ipv4/ip_forward
for n in $(find /proc/sys/net/ipv4 -name rp_filter) ; do
  echo "0" > $n
done
</snip>
```

Running tcpdump on both eth0 and ppp0, I saw that e.g. a http request from one of the local machines is actually passed via ppp0 to the remote host. However, all reply packets from that box are never passed back to eth0, so this looks to me as if masquerading somehow fails. I got one comment from the iptables forum that tzh the third message on the ppp0 side means the a tcp reset has been sent back by G4 upon the reception of the packet from the outside world:

```
[root@antares root]# tcpdump -nn -i eth0 tcp port 80
18:16:21.012143 IP 192.168.42.4.49223 > 213.95.27.115.80: S 2685214081:2685214081(0) win 65535 <mss 1460,nop,wscale 0,nop,nop,timestamp 2148180757 0>
18:16:23.779283 IP 192.168.42.4.49223 > 213.95.27.115.80: S 2685214081:2685214081(0) win 65535 <mss 1460,nop,wscale 0,nop,nop,timestamp 2148180762 0>
18:16:26.626863 IP 192.168.42.4.49224 > 213.95.27.115.80: S 2390183934:2390183934(0) win 65535 <mss 1460,nop,wscale 0,nop,nop,timestamp 2148180768 0>
(more messages snipped)
```

```
root@antares root]# tcpdump -nn -i ppp0 tcp port 80
18:16:21.012206 IP 84.44.131.113.49223 > 213.95.27.115.80: S 2685214081:2685214081(0) win 65535 <mss 1452,nop,wscale 0,nop,nop,timestamp 2148180757 0>
18:16:21.085651 IP 213.95.27.115.80 > 84.44.131.113.49223: S 2677460604:2677460604(0) ack 2685214082 win 5792 <mss 1460,nop,nop,timestamp 1472713132 2148180757,nop,wscale 2>
18:16:21.085748 IP 84.44.131.113.49223 > 213.95.27.115.80: R 2685214082:2685214082(0) win 0
18:16:23.779332 IP 84.44.131.113.49223 > 213.95.27.115.80: S 2685214081:2685214081(0) win 65535 <mss 1452,nop,wscale 0,nop,nop,timestamp 2148180762 0>
18:16:23.841268 IP 213.95.27.115.80 > 84.44.131.113.49223: S 2680216981:2680216981(0) ack 2685214082 win 5792 <mss 1460,nop,nop,timestamp 1472715888 2148180762,nop,wscale 2>
18:16:23.841326 IP 84.44.131.113.49223 > 213.95.27.115.80: R 2685214082:2685214082(0) win 0
(more messages snipped)
```

Running the same http access from the G4 *does* work without problems!

Does anyone know what I missed here? The same iptables setup (actually a lot stricter, i.e. a "real" firewall) worked fine with ISDN/ipp0. I tried various combinations of (below /proc/sys/net/ipv4/) rp_filter set to 0 and 1, as well as setting ip_dynaddr to 0, 1 or 2 – all without success.

I also verified that this setup is at least technically working; running the G4 under MacOS 10.3.9 client, with a little ipfw and natd fiddling the machine is doing full nat and firewalling as expected. However, as I

comp.os.linux.networking: HELP: NAT/Masquerading broken with 2.6.11 + pppoe (long)

usually use Linux, a running is really important for me.

HELP – I am really lost here, so any help/pointer would be really welcome!
As I am no kernel/networking guru, I also don't know where/how I could add more debugging to track down the problem. Maybe this is a PowerPC (i.e. endianness) specific problem? I meanwhile asked this question in a number of lists (inter alia netfilter, linux-net and the german ADSL forum), but did not get any helpful reply...

Thanks in advance,
Albrecht.

--

~~~~~  
Albrecht Dreß - Johanna-Kirchner-Straße 13 - D-53123 Bonn (Germany)  
Phone (+49) 228 6199571 - <mailto:albrecht.dress@arcor.de>  
GnuPG public key: <http://home.arcor.de/dralbrecht.dress/pubkey.asc>  
~~~~~