

netfilter: port restricted NAT

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-10/5291.html>

From: Alfred E. Heggstad (*alfredh_at_sxdesign.com*)

Date: 10/25/03

To: linux kernel <linux-kernel@vger.kernel.org>

Date: Sat, 25 Oct 2003 15:18:34 +0200

Hi

[this question was sent to netfilter mailing list but dropped]

I have a question about restricted NAT vs. port restricted NAT in the implementation of netfilter in 2.4.18 kernel. We are developing Voice-over-IP products using SIP for signalling, where the packets normally go over UDP port 5060.

For VoIP devices located behind NATs we use a protocol called "STUN – Simple Traversal of UDP Through NATs" (RFC3489) to probe for public interface ip/ports, and refresh connection tracking in netfilter for outgoing/incoming RTP streams. The SIP proxies and/or devices are located on the public network.

This works quite well with most SIP proxies but we have seen some cases where the response is "lost" in the gateway (linux 2.4.18) This is the case where the source port of the response is not the same as destination port of the request. I have looked around in these files:


```
net/ipv4/netfilter/ip_conntrack_proto_udp.c
net/ipv4/netfilter/ip_nat_proto_udp.c
```

and it looks like the connection tracker requires the response to come from the same port as the outgoing request was sent to.

Typical scenario:

[DEVICE 10.47.11.109]

```
|
|
V
src: 10.47.11.109:5060
dst: 80.80.80.80.5060
```



The packet is lost in the gateway.

>From RFC3489:

Restricted Cone: A restricted cone NAT is one where all requests from the same internal IP address and port are mapped to the same external IP address and port. Unlike a full cone NAT, an external host (with IP address X) can send a packet to the internal host only if the internal host had previously sent a packet to IP address X.

Port Restricted Cone: A port restricted cone NAT is like a restricted cone NAT, but the restriction includes port numbers. Specifically, an external host can send a packet, with source IP address X and source port P, to the internal host only if the internal host had previously sent a packet to IP address X and port P.

This behaviour I suspect would be close to the "Port Restricted Cone" definition. My main question is: With netfilter, is it possible at all to have only "Restricted Cone" with no source port checking and if yes how is this possible to configure?

I have read through iptables man page and searched www but could not find any reference to my problem.

For reference, here is the iptables NAT table on my gateway:

```
root@uranus:~# iptables -t nat -vL
Chain PREROUTING (policy ACCEPT 1762K packets, 285M bytes)
 pkts bytes target prot opt in out source
destination
```

Linux–Kernel: netfilter: port restricted NAT

Chain POSTROUTING (policy ACCEPT 833K packets, 209M bytes)

pkts bytes target prot opt in out source

destination

177K 11M MASQUERADE all -- any eth0 10.47.10.0/24

anywhere

107K 13M SNAT all -- any eth0 anywhere

anywhere to:213.187.186.100

Chain OUTPUT (policy ACCEPT 227K packets, 26M bytes)

pkts bytes target prot opt in out source

destination

Hopefully you are able to understand my question, thanks for any help.

/alfred

–

To unsubscribe from this list: send the line "unsubscribe linux–kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo–info.html>

Please read the FAQ at <http://www.tux.org/lkml/>