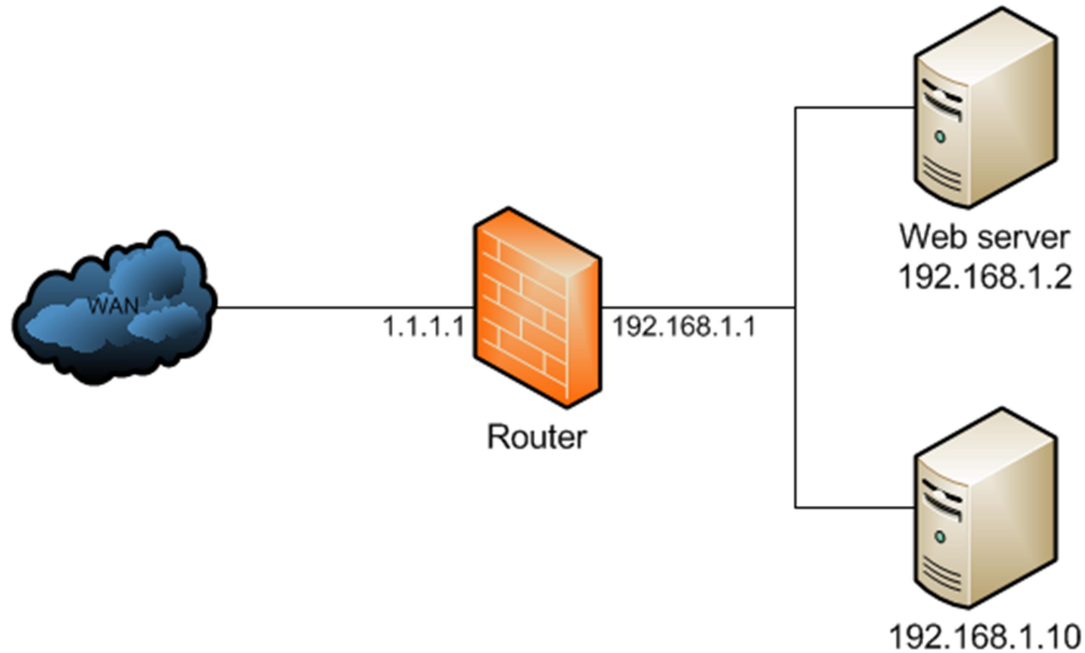# Hairpin NAT

In the below network topology a web server behind a router is on private IP address space, and the router performs NAT to forward traffic to its public IP address to the web server behind it.

The NAT configuration would look like below:

```
/ip firewall nat
add chain=dstnat dst-address=1.1.1.1 protocol=tcp dst-port=80 \
   action=dst-nat to-address=192.168.1.2
add chain=srcnat out-interface=WAN action=masquerade
```

When a client out on the Internet with IP address 2.2.2.2 establishes a connection to the web server, the router performs NAT as configured.

	Source IP 2.2.2.2	Destination IP 1.1.1.1
Step 1		

	Source IP 2.2.2.2	Destination IP 192.168.1.2
Step 2		

	Source IP 192.168.1.2	Destination IP 2.2.2.2
Step 3		

	Source IP 1.1.1.1	Destination IP 2.2.2.2
Step 4		

1. the client sends a packet with a source IP address of 2.2.2.2 to a destination IP address of 1.1.1.1 on port tcp/80 to request some web resource.
2. the router destination NATs the packet to 192.168.1.2 and replaces the destination IP address in the packet accordingly. The source IP address stays the same: 2.2.2.2.
3. the server replies to the client's request and the reply packet has a source IP address of 192.168.1.2 and a destination IP address of 2.2.2.2.
4. the router determines that the packet is part of a previous connection and undoes the destination NAT, and puts the original destination IP address into the source IP address field. The destination IP address is 2.2.2.2, and the source IP address is 1.1.1.1.

The client receives the reply packet it expects, and the connection is established.

When a client on the same internal network as the web server requests a connection to the web server's public IP address, the connection breaks.
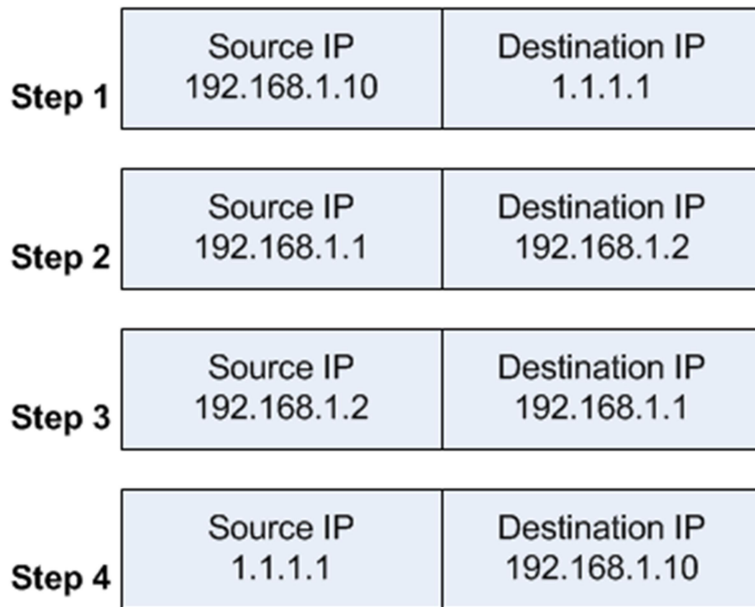
	Source IP 192.168.1.10	Destination IP 1.1.1.1
Step 1		

	Source IP 192.168.1.10	Destination IP 192.168.1.2
Step 2		

	Source IP 192.168.1.2	Destination IP 192.168.1.10
Step 3		

1. The client sends a packet with a source IP address of 192.168.1.10 to a destination IP address of 1.1.1.1 on port tcp/80 to request some web resource.

2. The router destination NATs the packet to 192.168.1.2 and replaces the destination IP address in the packet accordingly. The source IP address stays the same: 192.168.1.10.

3. The server replies to the client's request. However, the source IP address of the request is on the same subnet as the web server. The web server does not send the reply back to the router, but sends it back directly to 192.168.1.10 with a source IP address in the reply of 192.168.1.2.

The client receives the reply packet, but it discards it because it expects a packet back from 1.1.1.1, and not from 192.168.1.2. As far as the client is concerned the packet is invalid and not related to any connection the client previously attempted to establish.

To fix the issue, an additional NAT rule needs to be introduced on the router to enforce that all reply traffic flows through the router, despite the client and server being on the same subnet. The rule below is very specific to only apply to the traffic that the issue could occur with - if there are many servers the issue occurs with, the rule could be made broader to save having one such exception per forwarded service.

```
/ip firewall nat
add chain=srcnat src-address=192.168.1.0/24 \
  dst-address=192.168.1.2 protocol=tcp dst-port=80 \
  out-interface=LAN action=masquerade
```

	Source IP	Destination IP
Step 1	192.168.1.10	1.1.1.1

	Source IP	Destination IP
Step 2	192.168.1.1	192.168.1.2

	Source IP	Destination IP
Step 3	192.168.1.2	192.168.1.1

	Source IP	Destination IP
Step 4	1.1.1.1	192.168.1.10

With that additional rule, the flow now changes:

1. the client sends a packet with a source IP address of 192.168.1.10 to a destination IP address of 1.1.1.1 on port tcp/80 to request some web resource.

2. the router destination NATs the packet to 192.168.1.2 and replaces the destination IP address in the packet accordingly. It also source NATs the packet and replaces the source IP address in the packet with the IP address on its LAN interface. The destination IP address is 192.168.1.2, and the source IP address is 192.168.1.1.

3. the web server replies to the request and sends the reply with a source IP address of 192.168.1.2 back to the router's LAN interface IP address of 192.168.1.1.

4. the router determines that the packet is part of a previous connection and undoes both the source and destination NAT, and puts the original destination IP address of 1.1.1.1 into the source IP address field, and the original source IP address of 192.168.1.10 into the destination IP address field.

The client receives the reply packet it expects, and the connection is established.

However, the web server only ever sees a source IP address of 192.168.1.1 for all requests from internal clients regardless of the internal client's real IP address. There is no way to avoid this without either using a router that can do application level DNS inspection and can rewrite A records accordingly, or a split DNS server that serves the internal clients the internal server IP address and external clients the external server IP address.

This is called - among other terms - hair pin NAT because the traffic flow has clients enter the router through the same interface it leaves through, which when drawn looks like a hair pin.

-----------------o0o------------------