Snapt HAProxy TPROXY Manual

Version 1.0

Contents

Chapter 1: What is TPROXY 3
Chapter 2: How do I compile TPROXY support? 4
Chapter 3: How do I configure TPROXY? 5
Chapter 4: How do I deploy TPROXY? 6
Chapter 5: Final Steps 7
Chapter 6: Tips & Tricks 8

Chapter 1: What is TPROXY

Let's start with an explanation of what TPROXY is. You may have heard of it being used in load balancing, or in caching (with Squid). A problem that exists with almost all solutions that "proxy" traffic in some way is that the destination servers no longer see the origin clients IP address.

In the HAProxy example, let's say you have 3 email servers that you are load balancing. They see all email as coming from your HAProxy balancer, because technically it does, as it is proxying the inbound connections in order to load balance them. This is not on its own a problem, however, some anti-spam techniques use the origin IP in their checking (for example, RBLs). With caching it means you can't see the user who was browsing the site, only the proxies IP.

This is where TPROXY comes in, originally a kernel patch for Linux, it allows you to "spoof" the clients IP address, as if the load balancer was not there. When HAProxy sends the request to your mail server it looks as though it comes from the client that connected to HAProxy.

Chapter 2: How do I compile TPROXY support?

TPROXY is now included in the Linux kernel, so the only software modifications that are required (potentially) is for you to compile HAProxy with TPROXY support. You can get this from us directly if you enquire with us on the client site, or you can compile your own.

Download the latest HAProxy, extract it, and compile it like so --

```
make TARGET=linux26 USE_LINUX_TPROXY=1
make install target=linux26
```

This will typically put a haproxy binary in /usr/local/sbin/haproxy - some Linux distributions put it in /usr/sbin/haproxy by default, so if you had an existing haproxy you may wish to run the following command:

```
cp /usr/local/sbin/haproxy /usr/sbin/haproxy
```

Chapter 3: How do I configure TPROXY?

Setting a group to use TPROXY spoofing is quite easy in HAProxy, you need to add a single line to your group:

```
source 0.0.0.0 usesrc clientip
```

This tells HAProxy to use the client IP address as the source for all connections to the group. The complicated part comes in with iptables, the linux firewall system which takes care of the actual "spoofing" of the IP addresses.

In order to customize this, you should have a reasonably advanced understanding of the firewalling in Linux. Alternatively you can request that Snapt setup a configuration for you (with no charge), or assist you with yours, via the client site.

You require two pieces to begin the firewall side of this, the first is to add the iptables rules and the second is to mark the traffic for iptables. The easiest way to do this is to have the following script run at boot --

```
#!/bin/sh
/sbin/iptables -t mangle -N DIVERT
/sbin/iptables -t mangle -A PREROUTING -p tcp -m socket -j DIVERT
/sbin/iptables -t mangle -A DIVERT -j MARK --set-mark 1
/sbin/iptables -t mangle -A DIVERT -j ACCEPT

/sbin/ip rule add fwmark 1 lookup 100
/sbin/ip route add local 0.0.0.0/0 dev lo table 100
```

Chapter 4: How do I deploy TPROXY?

One complication that is often forgotten when using TPROXY is that the return traffic must pass through your HAProxy unit, because your servers will believe it has come from the client and not from the HAProxy unit (due to the IP spoofing). There are two easy ways of doing this; the first is if you have your HAProxy device act as a bridge in front of the traffic to your servers. This requires two network cards that are configured to bridge all traffic.

The second is to have a separate subnet that your servers are in, and have your HAProxy device act as a gateway on that subnet - ensuring all their traffic will go to HAProxy.

In this setup, you will also need to change your systems forwarding options (this may already be done depending on your installation) by running the following commands:

```
echo 1 > /proc/sys/net/ipv4/conf/all/forwarding
echo 2 > /proc/sys/net/ipv4/conf/default/rp_filter
echo 2 > /proc/sys/net/ipv4/conf/all/rp_filter
echo 0 > /proc/sys/net/ipv4/conf/eth0/rp_filter
```

To make this permanent you can add the following lines to /etc/sysctl.conf --

```
net.ipv4.forwarding = 1
net.ipv4.conf.default.rp_filter = 2
net.ipv4.conf.all.rp_filter = 2
net.ipv4.conf.eth0.rp_filter = 0
```

Chapter 5: Final Steps

All that's left to do now is to test it! Remember the following most important pieces:

1. Make sure all the commands are running at boot as well, for when you reboot your system.

2. The traffic from your web servers MUST go to the HAProxy device.

3. "source 0.0.0.0 usesrc clientip" tells a group to spoof.

Also remember that we provide full support, and you can contact us at any time for help with this setup.

Chapter 6: Tips & Tricks

Getting a weird error about privelages? TPROXY requires haproxy runs as root, so remove any user, group, uid, and gid options from your configuration.

Want an easy place to add the iptables rules shown above? Try /etc/rc.local