Control BGP from your Applications

RIPE 63 – November 2011

Thomas Mangin
Exa Networks

exa networks

Whatever a speaker is missing in depth he will compensate for in length

Montesquieu

Why ?

Dynamically Change Routing

DDOS filtering (RTBH)
IWF interception

Traffic engineering
Suspend customers
...

IP announcement ..

AnyCast
Control “cloud” IPs from a central location
Active / Passive servers solution with service monitoring

exa networks

How ?

Permanent configuration generation

1 - Regenerating BIRD/Quagga/OpenBGPD configuration on change
2 - Getting the daemon to reload its configuration
3 - Go back to 1

There must be a better way ...

OpenBGPD bgpctl
BIRD birdc
Quagga / Zebra telnet ..

There must be a better way .....

Logic will get you from A to B. Imagination will take you everywhere
Albert Einstein

Demonstration !

flap.sh

1 - take your favourite language : perl, python, lua, C, shell, french ! ...
2 - create a forever loop
3 - print what you want to do ...
4 - ... profit ?

```
#!/bin/sh

# ignore Control C
trap '' SIGINT

while `true`;
do
    echo "announce route 192.0.2.1 next-hop 10.0.0.1"
    sleep 10
    echo "withdraw route 192.0.2.1 next-hop 10.0.0.1"
    sleep 10
done
```

Integration

It is that simple

BGP configuration

```
neighbor 192.168.127.128 {
        description "will flap a route until told otherwise";
        router-id 198.111.227.39;
        local-address 192.168.127.1;
        local-as 65533;
        peer-as 65533;

        # add and remove routes when flap.sh prints
        process loving-flaps {
                run etc/processes/flap.sh;
        }
}
```


Success is a result, not a goal

Flaubert

Want simpler !!

BGP configuration

```
neighbor 192.168.127.128 {
        router-id 198.111.227.39;
        local-address 192.168.127.1;
        local-as 65533;
        peer-as 65533;

        process default-name-for-watchdog {
             run etc/processes/monitor.sh;
        }

        static {
              route 172.10.0.0/16 next-hop 192.0.2.1 watchdog service-one;
        }
}
```

Want Simpler ?

The watchdog ...

```
#!/bin/sh
trap '' SIGINT
while `true`;
do
    state=`check-if-all-ok`
        if [ “$state” = “up” ]; then
            echo "announce watchdog service-one"
        fi
        if [ “$state” = “down” ]; then
            echo "withdraw watchdog service-one”
        fi
        # pick its name from the process section name
        echo “announce watchdog”
        sleep 5
done
```

Flow Routes

Use BGP to transmit firewall like rules

- RFC 5575, Juniper routers only (atm)
- Can be used to transproxy in the core

Match possible components making the flow

- Prefix (source and destination)
- IP Protocol (list of <action, value>)
- Port (source, destination, either)
- ICMP (type, code), TCP flag, Packet Len, DSCP value
- Fragment (don't, is, first, last)

Then take action

- Drop, Rate-limit, Redirect

exabpg is the only OSS application to support Flow Routes

Be aware of line rate limitations when sending Flow Specs - test in a lab first.

The secret of business is to know something that nobody else knows

Aristotle Onassis

Example

```
neighbor 82.219.4.254 {
        description "Juniper router";
        router-id 10.0.0.1;
        local-address 10.0.0.1;
        local-as 65500;
        peer-as 65533;
        graceful-restart 5;

        flow {
                route optional-name-of-the-route {
                        match {
                                source 10.0.0.1/32;
                                source 10.0.0.9/32;
                                destination 192.168.0.1/32;
                                # port =80 =8080;
                                # destination-port >8080&<8088 =3128;
                                # source-port >1024;
                                # protocol [ tcp udp ];
                                # protocol tcp;
                                # packet-length >200&<300 >400&<500;
                                # fragment not-a-fragment;
                                # fragment [ first-fragment last-fragment ];
                                # icmp-type [ unreachable echo-request echo-reply ];
                                # icmp-code [ host-unreachable network-unreachable ];
                                # tcp-flags [ urgent rst ];
                                # dscp [ 10 20 ];

                        }
                        then {
                                # discard;
                                # rate-limit 9600;
                                # redirect 1.2.3.4:5678;
                                redirect 65500:12345;
                                community [30740:0 30740:30740];
                        }
                }
        }
}
```

exa networks

Get it ...

http://code.google.com/p/exabpg/

apt-get install exabpg

Questions ?

<first>@<last>.com

exa networks

Judge a man by his questions rather than by his answers

Voltaire