

# BGP Route Injection

**Thomas Mangin**, Exa Networks

[thomas.mangin@exa-networks.co.uk](mailto:thomas.mangin@exa-networks.co.uk)

**Andy Davidson**, NetSumo

[andy.davidson@netsumo.com](mailto:andy.davidson@netsumo.com)

LINUX69 ~ 15<sup>th</sup> February 2010


# Agenda

- Route Injection concepts
- Flow Spec Concepts
- Flow Spec tools


# Agenda

- **Route Injection concepts**
- Flow Spec Concepts
- Flow Spec tools

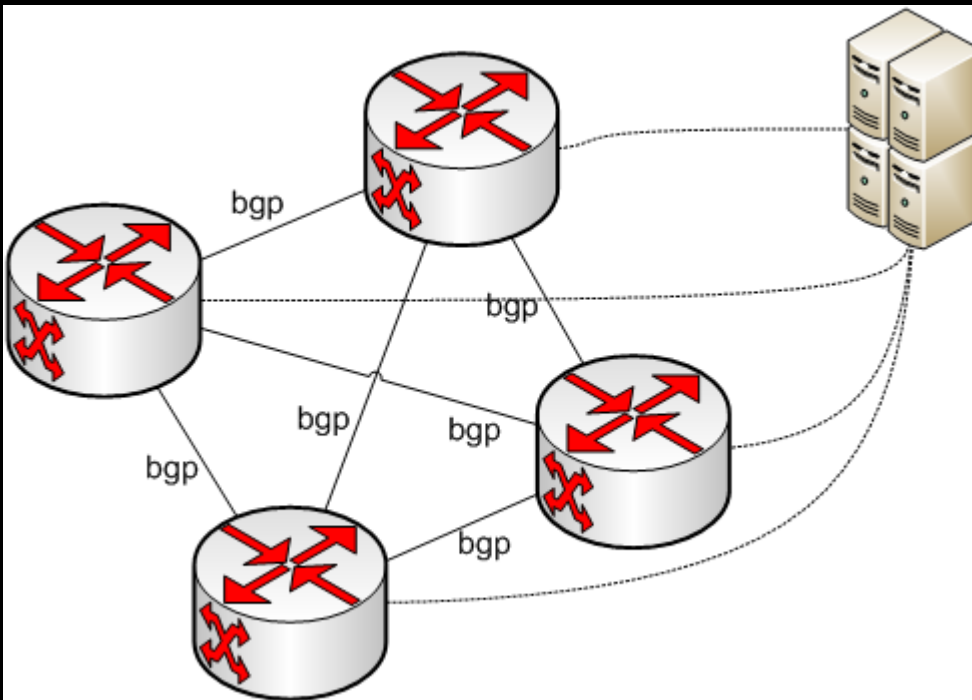
# Injection?



# Injection?




# Injection?



**Anycast**


# Injection?

Suspend  
Service



**Anycast**

# Injection?




**Suspend  
Service**

**Anycast**

**Compliance**

# Injection?




**Suspend  
Service**

**Migration**

**Anycast**

**Compliance**

# Injection?



**Suspend  
Service**

**Migration**

**Anycast**

**Compliance**

**Security / Flow Spec**

# EXA BGP Injector

- <http://bgp.exa.org.uk/>
- Simple, Juniper like configuration
- Injects prefixes with **community**, **next-hop**, and **local-preference** attributes
- iBGP or eBGP
- You may be doing this already with Quagga/BGPD/BIRD

# EXA BGP Injector Config

- neighbor 192.0.2.1 {  
    description "My router";  
    local-address 192.168.1.1;  
    local-as 65000;  
    peer-as 65000;  
    static {  
        route 10.0.1.0/24 next-hop 10.1.1.1 community  
65000:10 local-preference 999;  
    }  
}
- See <http://bgp.exa.org.uk/> for better examples

# EXA BGP Injector Notes

- Does not need full mesh peering
- Also allows injection of arbitrary AS-PATH  
route 10.0.1.0/24 next-hop 10.1.1.1 as-path [1 200000 3]
- Supports ASN32 and IPv6 (of course)
- BSD License – use and bend me ‘til I wear out.
- Flowscan and asn32 are new features (svn co)

# Agenda

- Route Injection concepts
- **Flow Spec Concepts**
- Flow Spec tools

# BGP Flow Specification info

- Use BGP to **distribute flow specification** patterns, and use routers to **filter traffic** matching these flows.

# BGP Flow Specification Method

- **New NLRI Address Family** for flow-spec.
- Express action, e.g. Accept, redirect, sample, discard using **extended communities**.
- Rules can be built against
  - Source and destination **prefix** and **ports**
  - **Protocol** (TCP, UDP, ICMP..)
  - ICMP **Type** and ICMP **Code**
  - **TCP Flags** (FIN, SYN, RST, PUSH, ACK, URGENT)
  - Packet **Length**
  - **DSCP**
  - **Fragmentation** (Don't frag, Fragmented, Last)

# BGP Flow Spec actions

- **Limit** traffic rate
- **Discard** traffic (traffic rate = 0)
- **Redirect** traffic into MPLS Tunnel
- You can alter these traffic behaviours without making a change to router configuration

# BGP Flow Spec advantages

- **Rapid rollout!** Rule distribution layer is already in place on your network..
- Finally use those expensive ASICs for something (filtering)
- Uses a **protocol that's well understood** by you and your NOC already

# BGP Flow Spec Standards

- New RFC – 5575 – August 2009
- Implemented by :
  - Juniper
  - Arbor
  - Exa BGP Injector
  - Maybe more .... ?
- Bug your Cisco Rep about this feature today 😊

# Example Uses

- **NOC** to stop Denial of Service. Ideally combined with Netflow tools.
- **Abuse Desks** to prevent spam from the outside or your own customers.
- **Enterprises** to enforce security policy
- **Compliance desks** to filter, e.g. IWF
- **BOFH**, to slow down boss's Counterstrike games!

# DOS Mitigation?

- Customer typically asks for **filter/ACL** to be applied on their ports
- Traffic must still **traverse the backbone**, not dropped at ingress/edge.
- Brutal filter allows **attack to succeed!**
- Attack traffic is more intelligent – TCP SYN flood / UDP fragment attack  
= **low PPS, high damage**

# Traditional response does not scale

- Lots of config generation/regeneration.
  - Slow to roll out
  - Risk of human error
  - Constant cat-and-mouse config race
- More automation = more time to **party**.

# Agenda

- Route Injection concepts
- Flow Spec Concepts
- **Flow Spec tools**

# Arbor

- You have just seen a complete presentation from Neotelecoms explaining the Arbor features, and how to configure your router.

# EXA BGP Example

```
static { ... }
  flow {
    route block-evil-spammer {
      match {
        source 195.66.232.40/32;
        port =25;
      }
      then {
        discard;
      }
    }
  }
}
```

# Putting it all together

- 1) **Identify attack vectors** with Netflow scripts and tools, e.g. Ntop / panoptis.
- 2) Automatically generate configuration and **filter specification** with Exa BGP. SIGHUP to propagate rules.
- 3) ...
- 4) **Profit**

# Future work

- Automation with more Netflow tools ???
- All vendors to support for Flowscan ???
- Inter provider flow-scan ???

Simple 😊

**Any Questions ?**

Any Answers ?

thomas.mangin@exa-networks.co.uk

andy.davidson@netsumo.com