BGP Flowspec

April 2008

Agenda

- The problem
- What is Flowspec?
- Components
- Validation
- What can we do with it?
- Junos Configuration

The problem

- Service Providers are being driven to detect and mitigate denial of service attacks destined towards key customers
  - Stop bad traffic from reaching customer
- Service Providers also want to
  - Stop bad traffic consuming resources on expensive transit links
  - Be able to position as a value add to customer

Layered solution

- CPE protection
  - Customer has UTM/DI/IDP
  - Granular inspection of every packet
- Provider upstream edge detection/blocking
  - Analysis of flow information
  - Dynamic filters applied to rate limit, block or redirect specific attack traffic
  - Eliminate human error or delay associated with traditional access list mitigation
- Centralised cleaning solution
  - Value add for customer that doesn't have deep inspection capability
  - Forensic analysis / packet capture

BGP Flowspec

- Use BGP to distribute flow specification filter and dynamically filter on routers
  - Introduced in Junos 7.2
  - New BGP NLRI address family
  - Use extended communities to specify action (accept, discard, rate-limit, sample, redirect)
  - Match on a combination of source/dest prefix, source/dest port, ICMP type/code, pack size, DSCP, TCP flag, fragment encoding etc.

What is BGP Flow-Spec

- RFC 5575 - Dissemination of Flow Specification Rules
- Defines a method for the originator of a BGP NLRI to define and advertise a flow filter to its peers via BGP.
- Multi vendor support
  - Co-authored with Cisco, Arbor, NTT/Verio
- Authors:
  - Jared Mauch
  - Danny McPherson
  - Robert Raszuk
  - Barry Greene
  - Pedro Marques
  - Nischal Sheth

What is BGP Flow-Spec

- Defines a way to carry “flow” in BGP
  - New Address family for BGP
    - NLRI type (afi=1, safi=133)
- Defines operations to perform on flows
  - Sends an “action” in a BGP Update
- Defines a Model for Validation

Address family identifier / sub address family indicator

Component Types

- T1 Destination Address
- T2 Source Address
- T3 IP Protocol
- T4 Port (source or dest)
- T5 Destination port
- T6 Source Port
- T7 ICMP type
- T8 ICMP code
- T9 TCP flags
- T10 Packet length
- T11 DSCP
- T12 Fragment Encoding

Actions

- Carried as extended BGP communities
- Type 0x8006 Traffic-Rate
- Type 0x8007 Traffic-Action
  - Bit 0 Action set to “action or not “ (filter or not)
  - Bit 1 Sample log the packets
- Type 0x8008 Redirect
  - Send traffic to another VRF for collection

Flow Validation

- Need to validate by default to prevent spoofing

- Rules

  a) The "originator" of a flow route matches the "originator" of the best match unicast route for the destination address that is embedded in the route.

  b) There are no more-specific unicast routes, when compared to destination address of the flow route, for which the active route has been received from a different next-hop autonomous-system.

Disabling Validation

- Validate against a policy

```
family inet {
    flow {
    no-validate <policy>; "Validation procedure is skipped
    for
                        routes that match this policy";
    }
}
```

What can we do with it

- Allows Customers to set their own firewalls on SP core.
  - Validation rules will avoid spoofing of flow NLRI
- Provides a tool for the NOC to quickly react to DDOS attacks.

Distributed DOS attack
In the “old” days

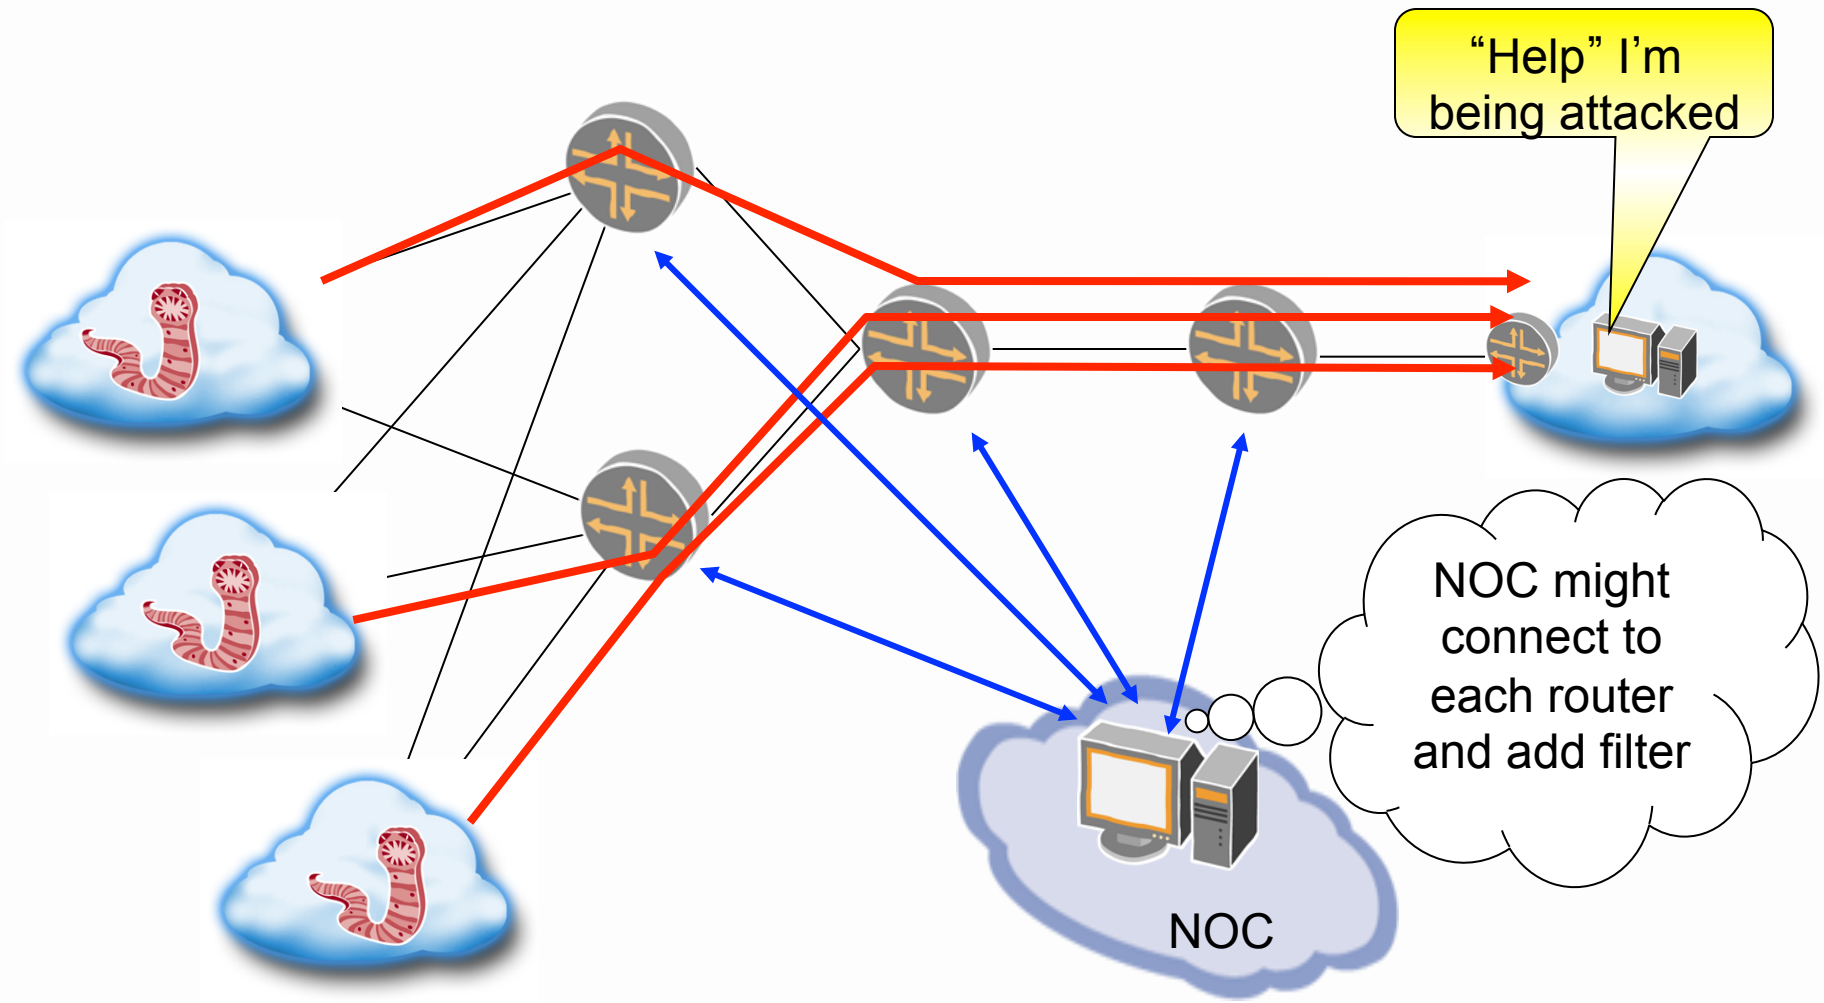

The General Concept – micro view

- CPE can now react to a DOS attack

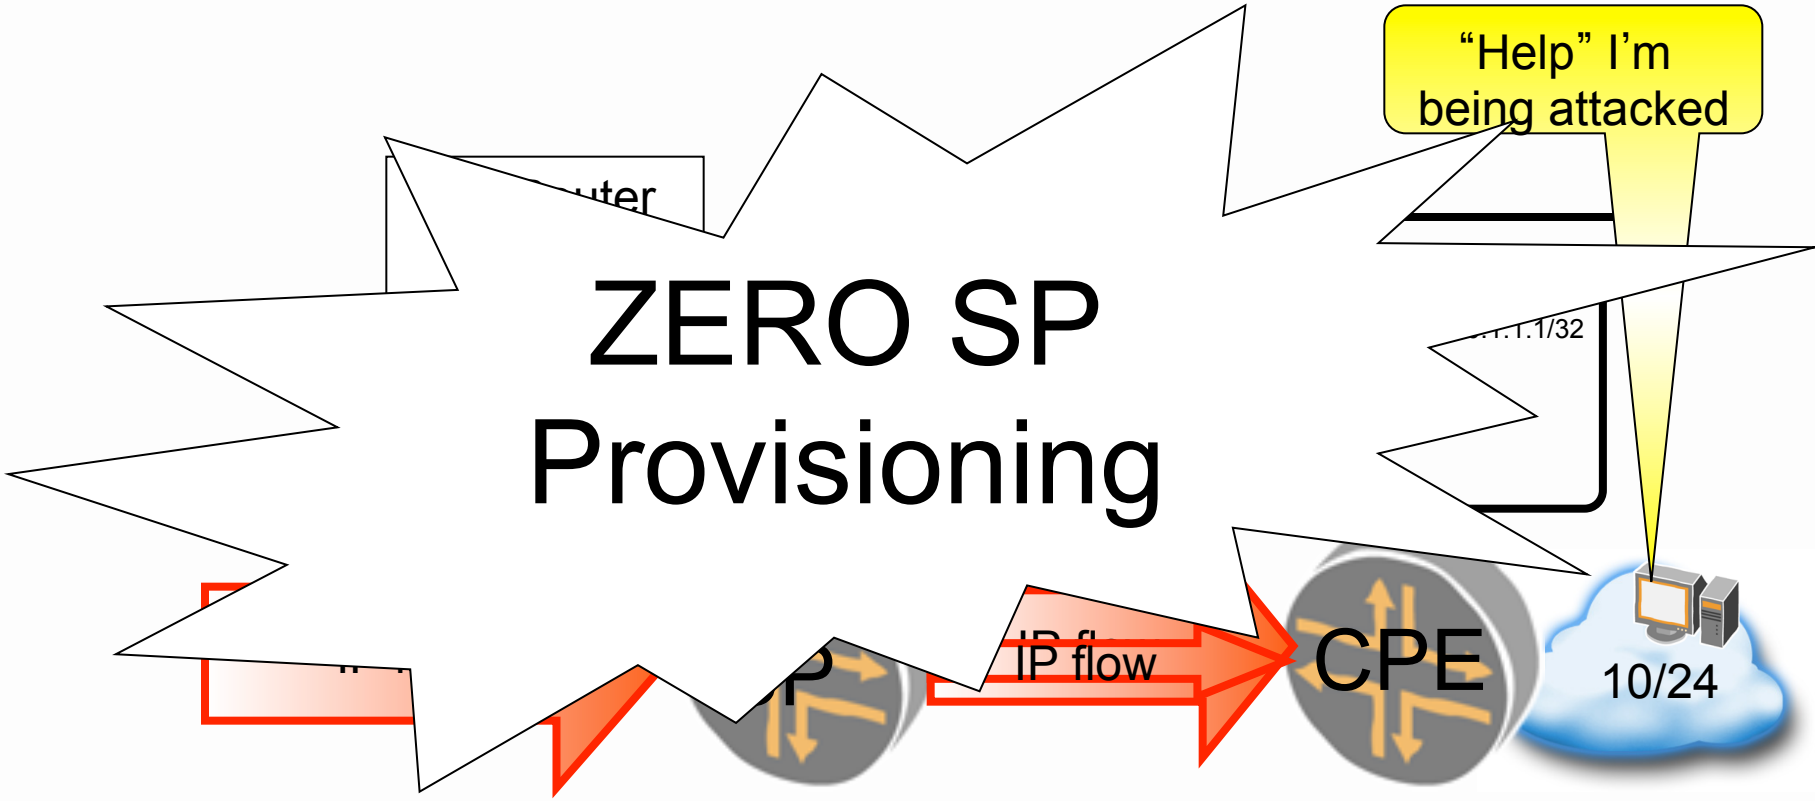

In model for monitoring, flow is small part of picture

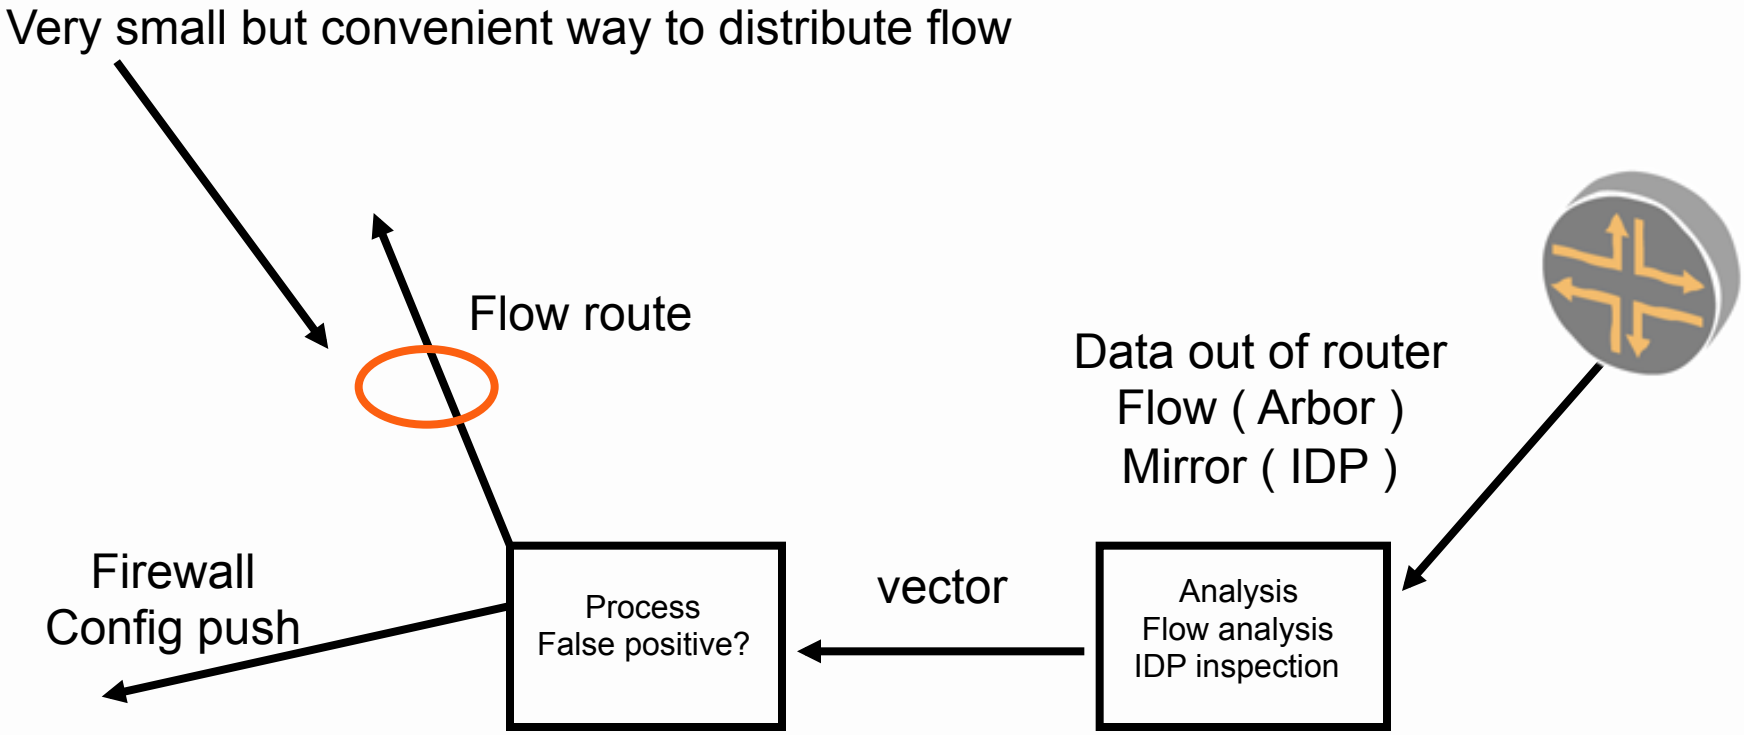

Distributed DOS attack
CPE Controlled

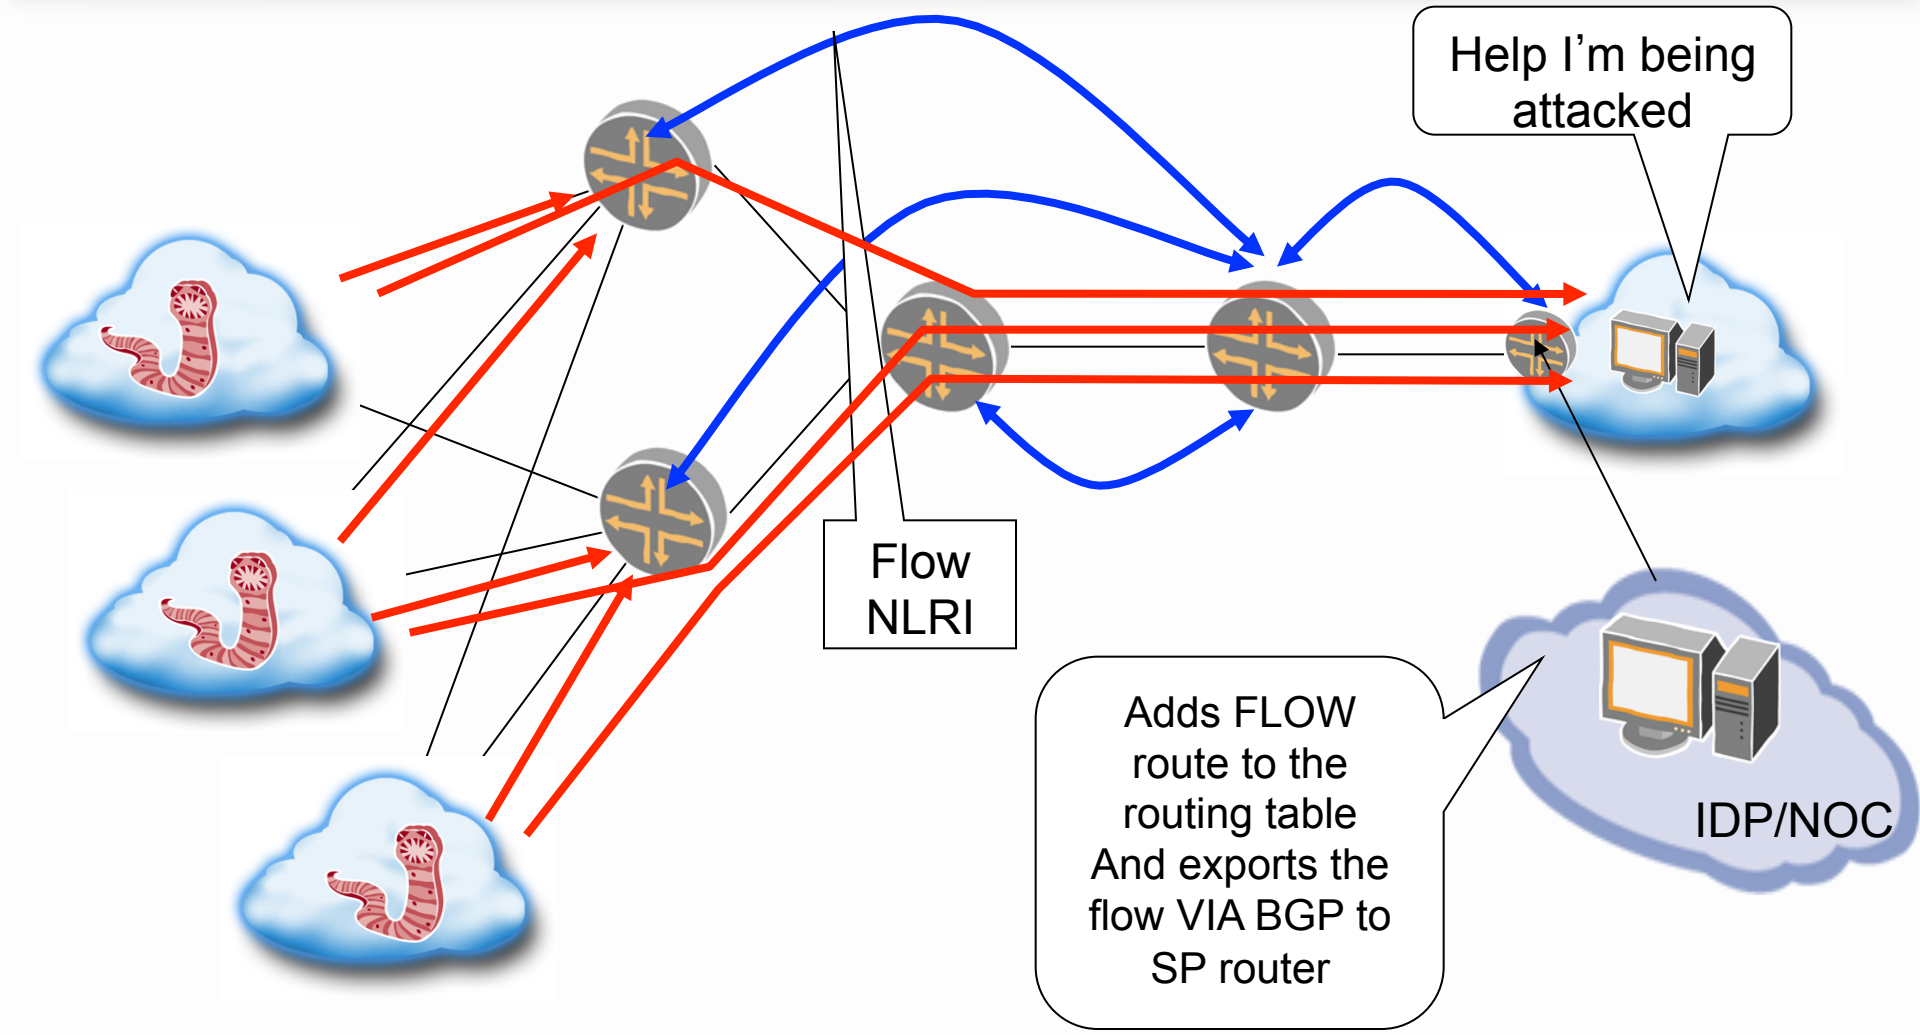

Comparisons with current filtering methods

- Many SP's already use prefix based filters
  - Match on community
  - Set next-hop discard
  - ONLY works for destination prefix
- Flow adds granularity to this
  - Match on components
    - SA / DA / Proto / length..
  - Don't have to discard
    - Rate limit
    - Sample
    - Forwarding-class

Configuration Options Define FLOW

```
routing-options {
    flow {
        route <name> {
            match {
                destination;
                source ;
                protocol ;
                port ;
                destination-port ;
                source-port ;
                icmp-code ;
                icmp-type ;
                tcp-flags ;
                packet-length ;
                dscp ;
                fragment [
                    dont-fragment
                    not-a-fragment
                    is-fragment
                    first-fragment
                    last-fragment
                ]
            }
```

```
            then {
                accept;
                discard;
                next-term;
                rate-limit;
                sample;
                routing-instance;
            }
        }
    }
}
```

```
[edit protocols bgp]
group <name> {
    family inet flow;

    neighbor <a.b.c.d> {
        family inet flow;
    }
}
```

Configuration Example Routing Options

- Define Flow routes

```
routing-options {
    flow {
        route filter {
            match destination 192.168.21.0/24;
            then {
                community test;
                rate-limit 32k;
            }
        }
    }
}
```

Configuration example BGP

- Add family flow to BGP peers

```
Protocols {
    bgp {
        group int {
            type internal;
            local-address 20.2.2.2;
              family inet {
                  unicast;
                  flow; <<<
              }
    neighbor 20.3.3.3;
}
```

Configuration example

- Define Non-Validation

```
show protocols bgp group int {
    type internal;
    local-address 20.3.3.3;
    family inet {
        unicast;
        flow {
            no-validate test;
        }
    }
    neighbor 20.2.2.2;
}
```

Diagnostics

- show route receive-protocol bgp
  - Shows received NLRI
- show route advertising-protocol bgp
  - Shows advertised NLRI
- show route flow
  - show active flow routes
- show route table inetflow.0
  - Shows actual defined flow routes (from routing options)
- show firewall
  - Shows installed flow filters and counters

Show Firewall

```
lab@Darstardly-re0# run show firewall


Counters:
Name                                                Bytes              Packets
192.168.21/24,*                                      28672              112
Policers:
Name                                                Packets
192.168.21/24,*                                      112

[edit]
lab@Darstardly-re0#
```

Who's using it

- Internet 2
- TimeWarner
- others looking into it
  - Dozens !

Big Motivation is VoIP

Common questions

- Spoofing
  - Validation will prevent this
- Why BGP
  - Its there
- What's stopped auto configuration efforts in the past?
  - AS boundaries
  - NO tools that work
    - Configure >100 routers in seconds “Danny McPherson”

Arbor BGP flowspec integration

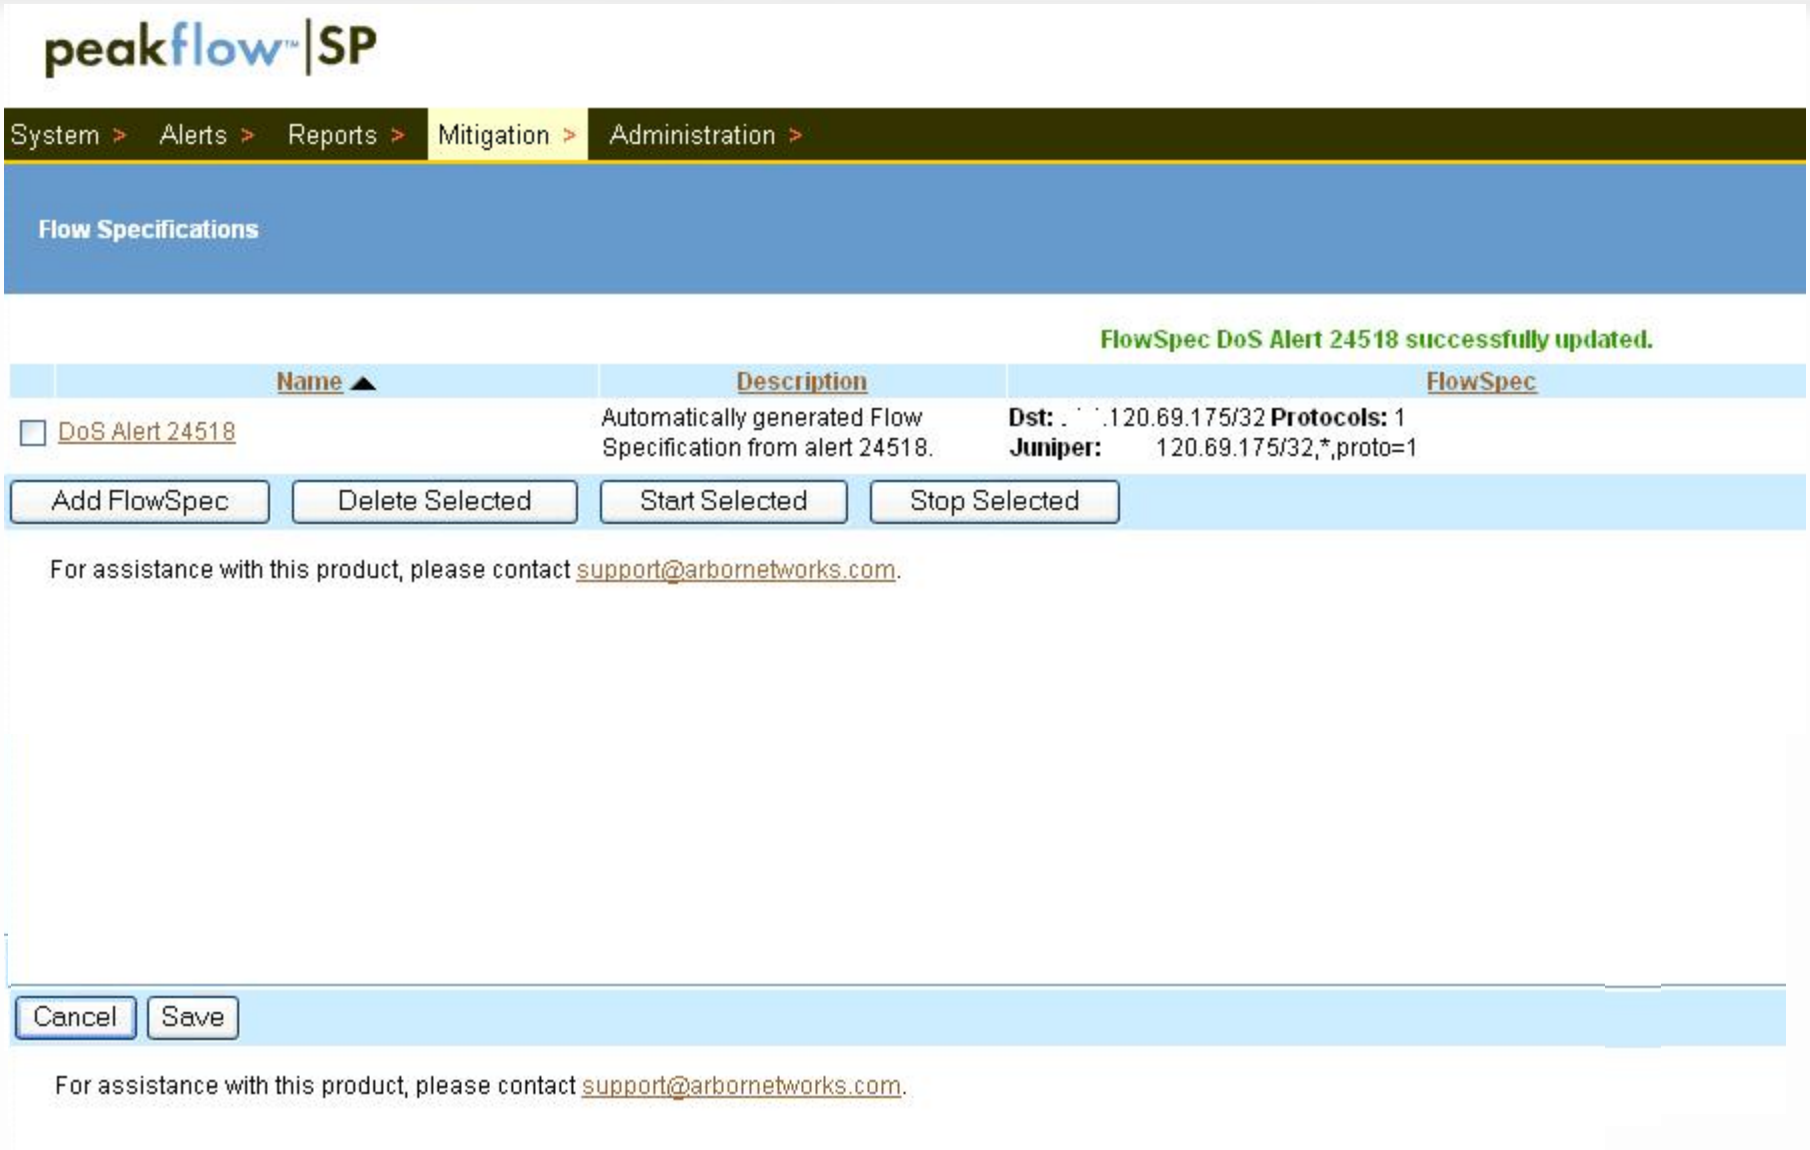

Things to think about…

- Propagation of filters to SP peers?
- Use in lawful intercept?

References

- http://www.nanog.org/mtg-0610/lozano.html
- http://tools.ietf.org/id/draft-marques-idr-flow-spec-04.txt
- http://www.ietf.org/proceedings/07jul/slides/idr-0.pdf