BGP Flowspec

Basics / DDoS Mitigations

Christoph Loibl

christoph.loibl@nextlayer.at

BGP Flowspec NLRI (RFC5575)

Rapidly deploy access control lists / flow-filters to routers
ie. during DDoS mitigation (not limited to that)

- RFC5575 defines a BGP NLRI to exchange flow specification rules via BGP
- Intra- and inter-AS distribution of flow-filter rules
- Minimal flow-filter verification based on unicast-routing announcement

Example – DDoS Mitigation

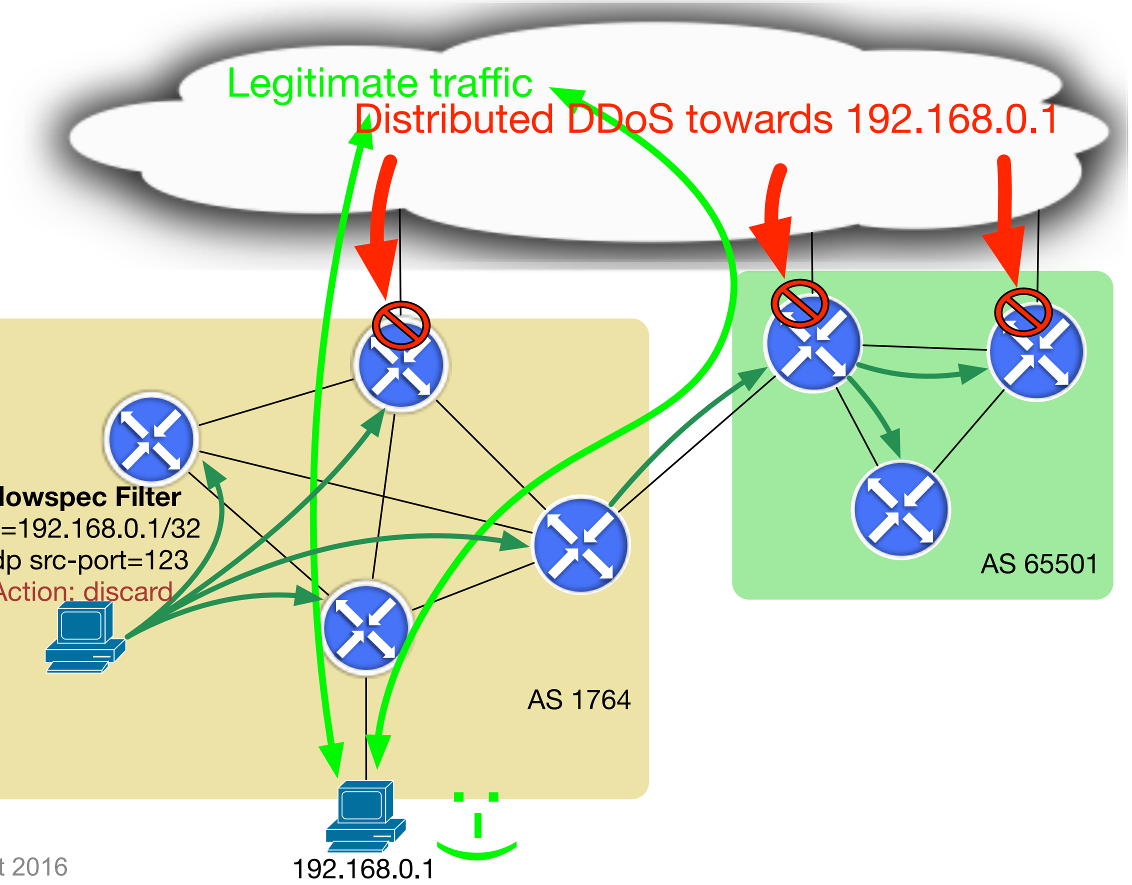

Flowspec Filters n-Tuples

Matching criterias (NLRI value)

- Source/Destination prefix
- IP protocol
- Source/Destination port
- ICMP type/code
- TCP flags
- Packet length
- DSCP
- Fragment

Filtering actions (community)

- Traffic rate
- Traffic action
- Redirect
- Traffic marking

Sorting Algorithm for Filter Rules

Why?

- Multiple filter rules with different actions may match a given packet
- Consistency over the network
- Time of arrival of the BGP announcement cannot be used for ordering

Order defined by the RFC5575:

- Sorting algorithm compares components of the matching criteria
- In general: More detailed criteria end up on top of the list

How to get started?

- No additional expensive hardware required!
- Current routing platforms support flowspec address-family
- Flowspec makes use of the existing BGP infrastructure
  (minimal configuration changes to the network required)

Originating Flowspec filters on demand via …

… configuration on a single router in the network (like static routes)

… a routing daemon like ExaBGP

Example Juniper Configuration

```
nextlayer@NL-LAB-AX1> show configuration protocols bgp group FLOWSPEC

local-address 10.0.0.1;
family inet {
    flow {
        no-validate ACCEPT;
    }
}
peer-as 65000;
neighbor 10.0.0.254;

{master:0}
```

Example ExaBGP Configuration

```
route {
   match {
      destination 192.168.0.1/32;
      source-port =1900 =19 =123 =53;
      protocol udp;
   }
   then { rate-limit 1250000; }
}
route {
   match {
      destination 192.168.0.1/32;
      destination-port =80 =443;
      protocol tcp;
   }
   then { redirect 1764:666; }
}
route {
   match {
      destination 192.168.0.1/32;
   }
   then { discard; }
}
```

Verifying the Received “Routes” 1/2

```
nextlayer@NL-LAB-AX1> show route table inetflow.0

inetflow.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.1,*/term:3
                   *[BGP/170] 00:04:26, localpref 100, from 10.0.0.254
                      AS path: I, validation-state: unverified
                      Fictitious
192.168.0.1,*,proto=6,dstport=80,=443/term:1
                   *[BGP/170] 00:04:26, localpref 100, from 10.0.0.254
                      AS path: I, validation-state: unverified
                      Fictitious
192.168.0.1,*,proto=17,srcport=1900,=19,=123,=53/term:2
                   *[BGP/170] 00:04:26, localpref 100, from 10.0.0.254
                      AS path: I, validation-state: unverified
                      Fictitious

{master:0}
```

Verifying the Received “Routes” 2/2

```
nextlayer@NL-LAB-AX1> show route table inetflow.0 detail

inetflow.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.0.1,*/term:3 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x9523604
                Next-hop reference count: 3
                State: <Active Int Ext>
                Local AS: 65000 Peer AS: 65000
                Age: 7:05
                Validation State: unverified
                Task: BGP_65000.10.0.0.254
                Announcement bits (1): 0-Flow
                AS path: I
                Communities: traffic-rate:0:0
                Accepted
                Localpref: 100
                Router ID: 10.0.0.254
```

Verifying the Resulting Filter

```
nextlayer@NL-LAB-AX1> show firewall filter __flowspec_default_inet__

Filter: __flowspec_default_inet__
Counters:
Name                                                 Bytes    Packets
192.168.0.1,*                                            0          0
192.168.0.1,*,proto=17,srcport=1900,=19,=123,=53         0          0
192.168.0.1,*,proto=6,dstport=80,=443                    0          0
Policers:
Name                                                 Bytes    Packets
10M_192.168.0.1,*,proto=17,srcport=1900,=19,=123,=53     0          0

{master:0}
```

The big picture

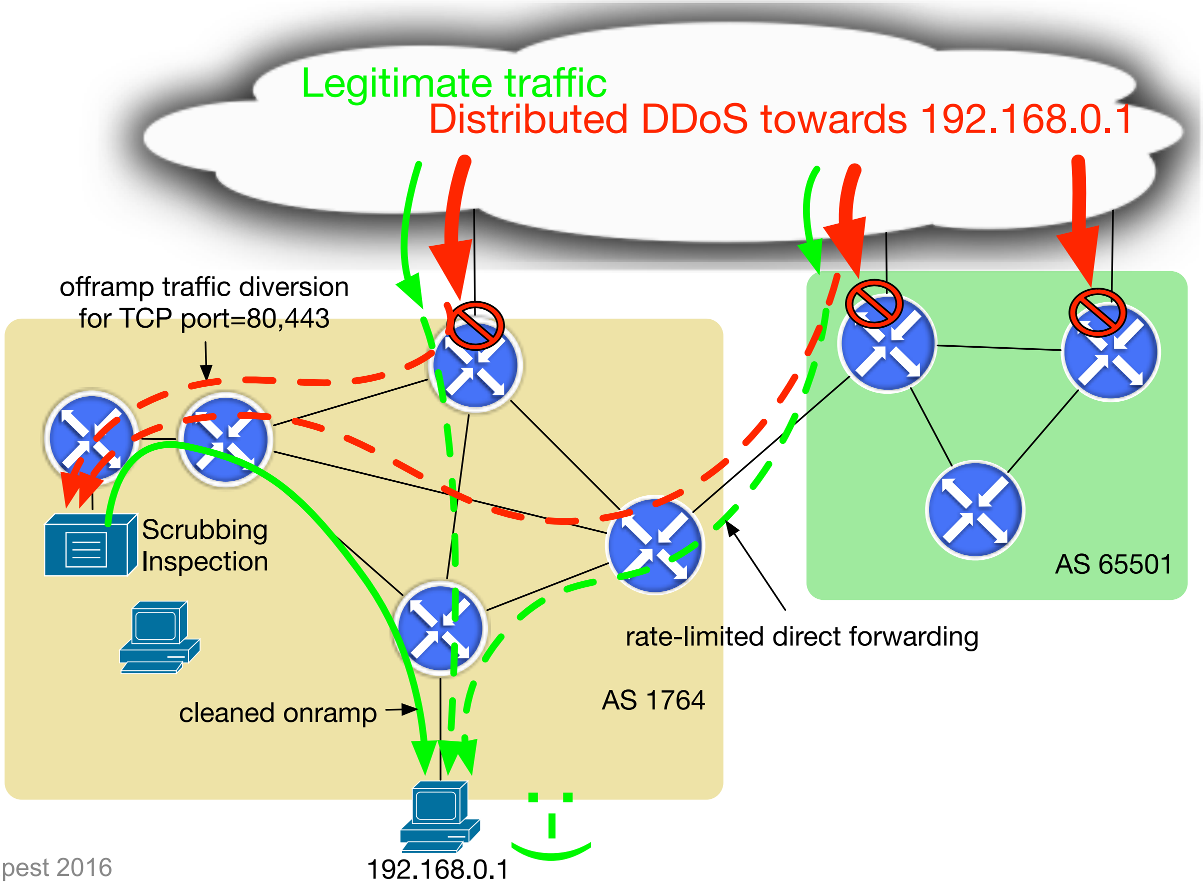

What next?

- Upcoming RFCs
  - draft-ietf-idr-flow-spec-v6

    Redefinition of RFC5575 for IPv6
  - draft-ietf-idr-flowspec-l2vpn

    Flowspec for L2VPN (ethernet frame filtering)
- Improvements / Industry adoption
  - More flexible vendor support
  - Operational experience / stability
  - Inter AS flowspec peerings / transit services?
  - IXP routeserver supporting flowspec?

Thank you.

Christoph Loibl

christoph.loibl@nextlayer.at