NETSCOUT | Arbor

Arm Yourself Against DDoS Attacks: Using BGP Flow Specification for Advanced Mitigation Architectures

IX Forum 11

Kleber Carriello, Sr. Consulting Engineer
Arbor Networks

Are you doing flowspec?

Flowspec is a lot like teenage sex, everyone has been talking about it, nobody really knows how to do it, everyone thinks everyone else is doing it, so everyone claims they are doing it.

Goals:

- Flowspec background;
- Primary use cases for flowspec – diverting traffic & blocking traffic;
- Keeping yourself safe when using Flowspec;

Customer 1

Brazil: Recent history of DDoS (2016 – 2017)

160 Million Packets Per Second Syn-Flood

week long +540Gbps gre-flood attack

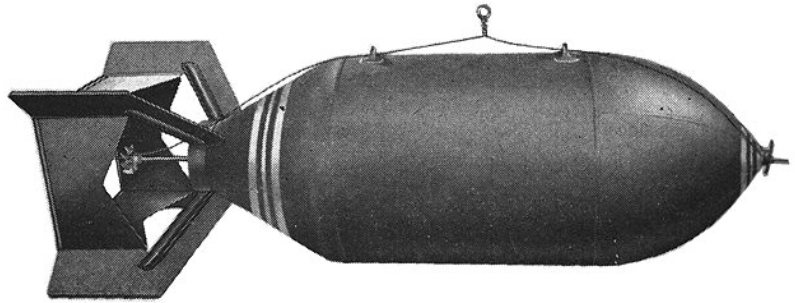

BGP Blackhole

S/RTBH

IDMS (TMS HD-1000)
Intelligent DDoS Mitigation System

BGP Flow Specification (Flowspec)

What is BGP Flow specification?

- Layer-4 Router ACLs that can be distributed and managed by BGP
- Provides for ability to match flows on the following items:
  - Source/Dest IP(s)
  - Source/Dest Ports
  - Protocol
  - Packet-Length*
  - TCP Flags*
  - Fragmentation Bits*
- Perform the following actions:
  - Rate-Limit BPS (0-drop)
  - Redirect-to-VRF
  - Set DSCP Values
  - Redirect to IP nexthop**

*Platform dependent

**RFC still draft & platform dependent

Manage distribution policy with BGP flexibility!

History of Flowspec

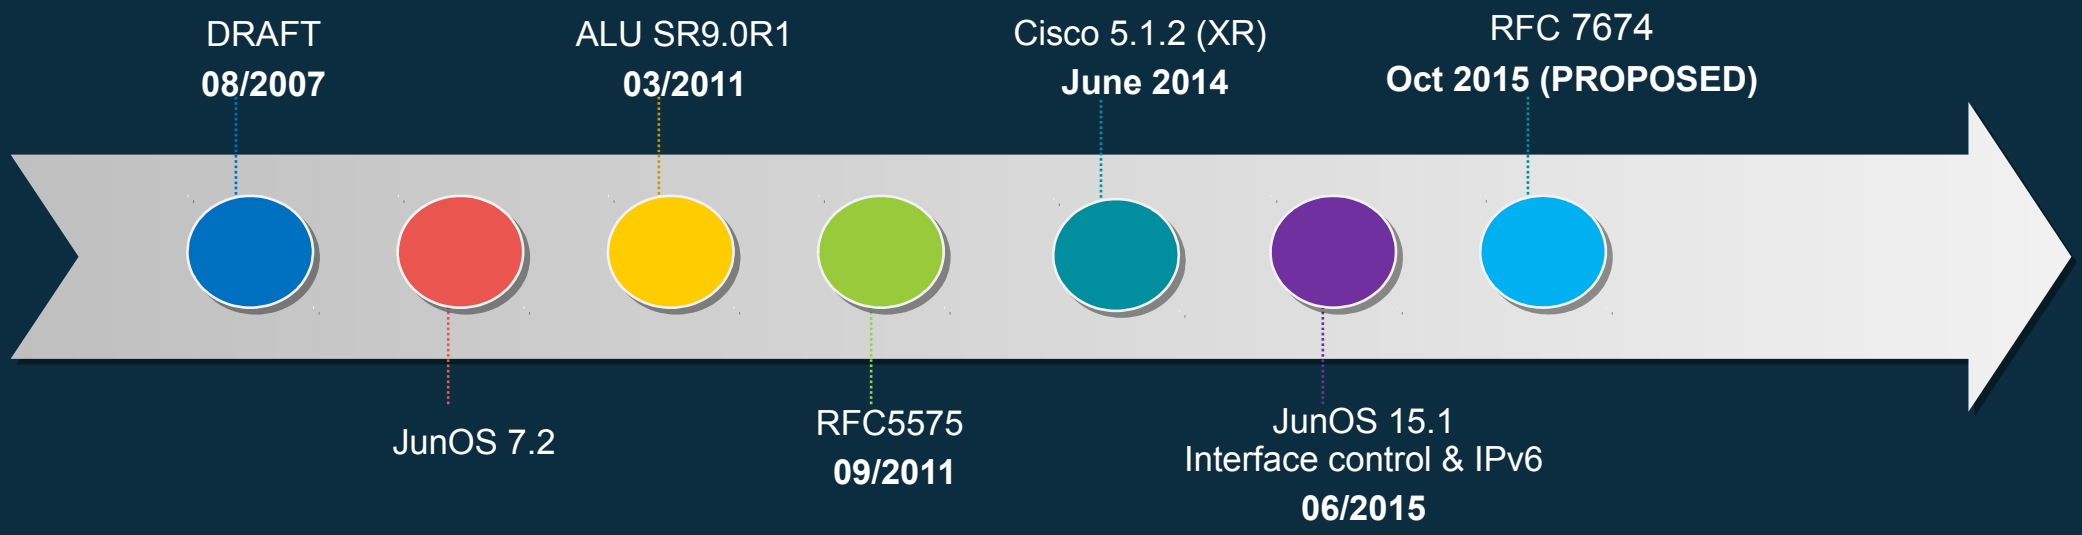

Challenges

- Vendor Support
- Specifications in flight
- Operational Challenges

Vendor Limits

- Feature parity approaching
- System Limits
- Encoding Differences

Vendor	Table-Limits
Alcatel-Lucent	512
Cisco	3000 (ASR9K)
Juniper	8000

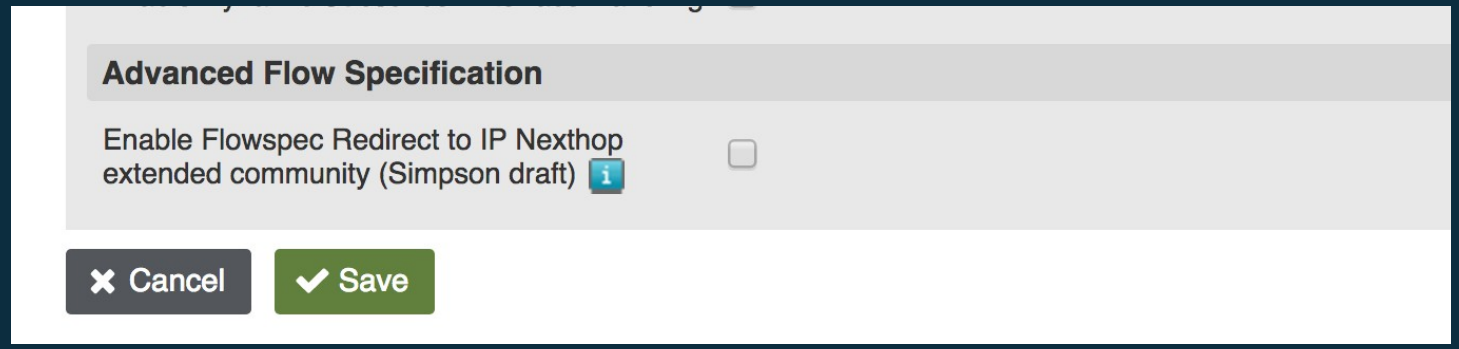

You need to understand your device's limits!

Flowspec Vendor Parity (Visibility)

- No standard for reporting
- Dropped traffic via Netflow
- Query router via API, SNMP, Yang
- Off-net solutions?

```
RP/0/RSP0/CPU0:edge-frankfurt#sh flowspec afi-all detail
Mon Mar 27 18:53:57.384 UTC

AFI: IPv4
  Flow           :Dest:6.6.6.6/32
    Actions      :Traffic-rate: 0 bps  (bgp.1)
    Statistics                       (packets/bytes)
      Matched              :             1536776/2332825968
      Dropped              :             1536776/2332825968
```

edge-frankfurt (4)	High	-	-	-	1500	654.2 Mbps 54.5 Kpps	554.3 Mbps 46.2 Kpps
Filtered by Router		OUT			1500	623.8 Mbps 52.0 Kpps	195.5 Mbps 16.3 Kpps
GigabitEthernet0/0/1/1 PEER:ESXI2-VMN2:AS65009:SANDERSYS		IN	Network	65009	1500	654.2 Mbps 54.5 Kpps	554.3 Mbps 46.2 Kpps
Bundle-Ether1 BB:CSJC-AS4.1004:AS65530		OUT			1500	334.2 Mbps 27.9 Kpps	168.3 Mbps 14.0 Kpps
Bundle-Ether2 BB:CNYK-AS4.1004:AS65530		OUT			1500	360.2 Mbps 30.0 Kpps	190.4 Mbps 15.9 Kpps

Flowspec Actions (traffic-rate)

- Traffic-rate with a value of ”0” means drop all traffic
- BPS only
- Limit per router
- Nothing guaranteed!

Per interface settings (Interface-set)

- Provides for policy application ingress on a router interface
- Essentially allows policy based routing (PBR)
- draft-litkowski-idr-flowspec-interfaceset-03.txt
- Specifies interfaces Flowspec rules are applied on the router

Benefits:

- ✓ Allows FS rules to only be applied to untrusted places on the network (Where your attack traffic comes from)
- ✓ Removes return-traffic complexities with scrubbing centers (No GRE!)
- ✓ Simplifies mitigation of East > West or Customer > Customer attacks

Flowspec – per interface control

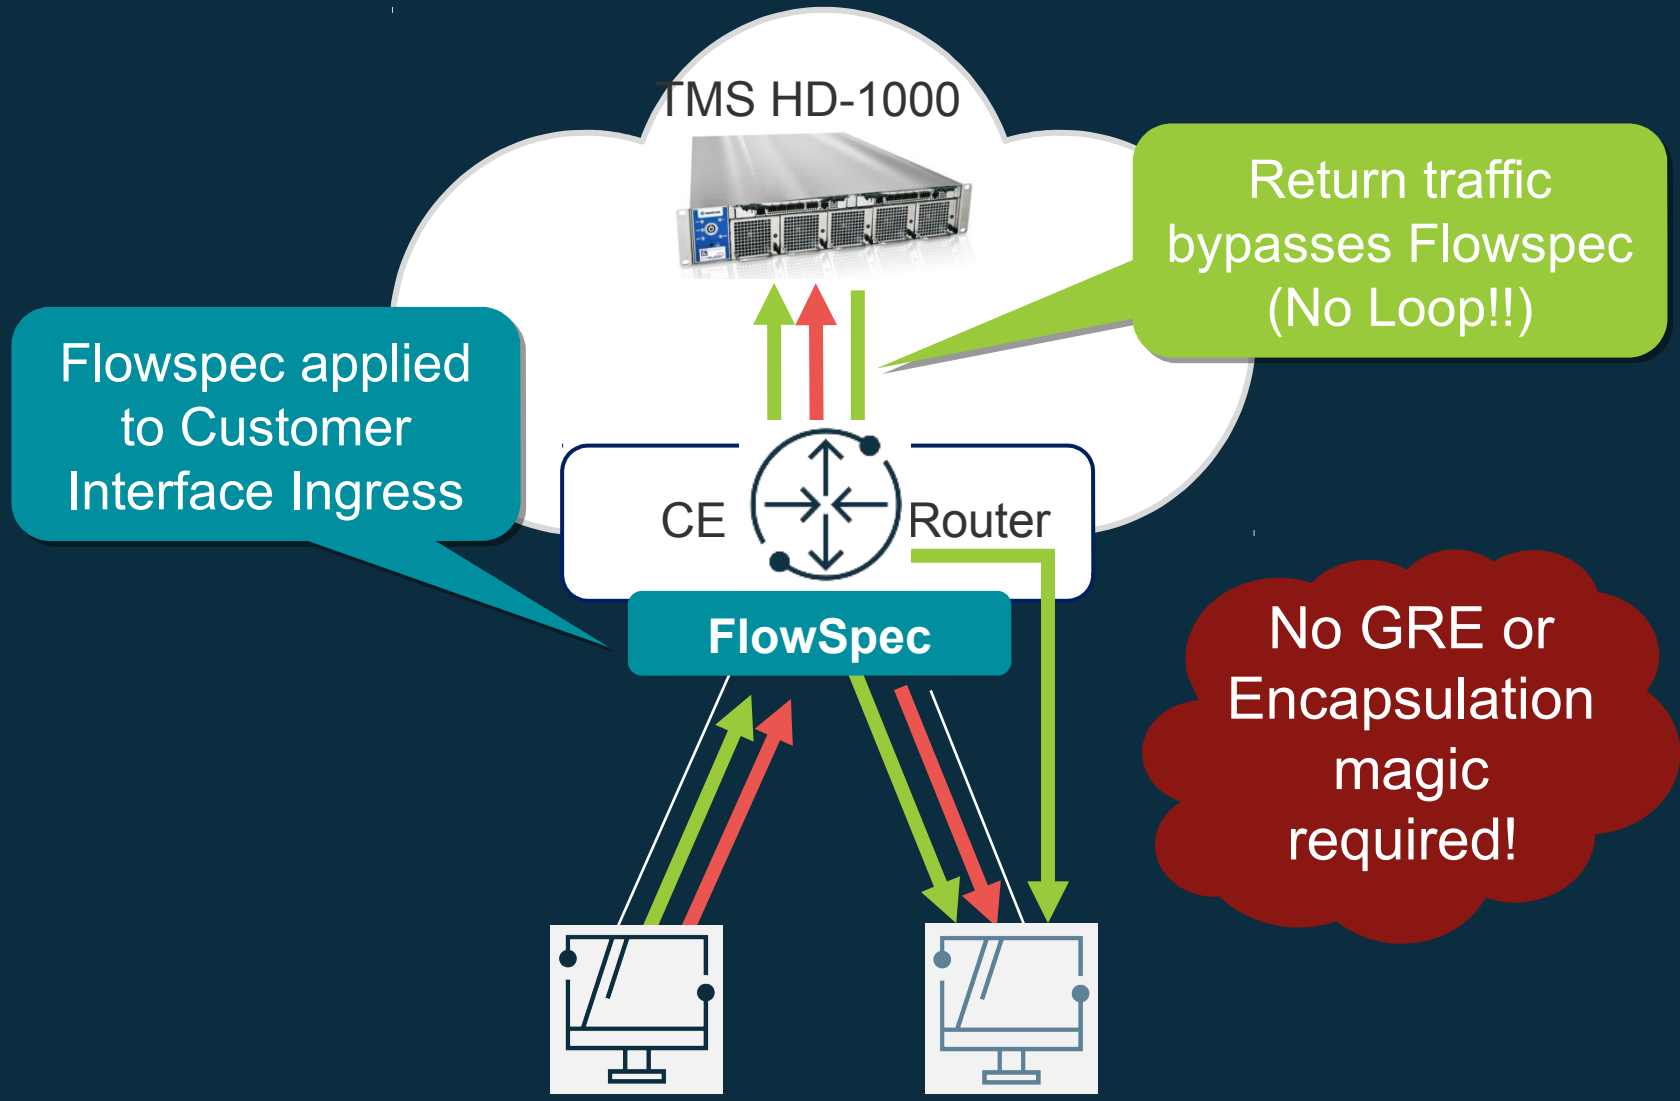

Enabling BGP Flow specification

- Enable the Flowspec address family
- Separate configuration for IPv4 & IPv6

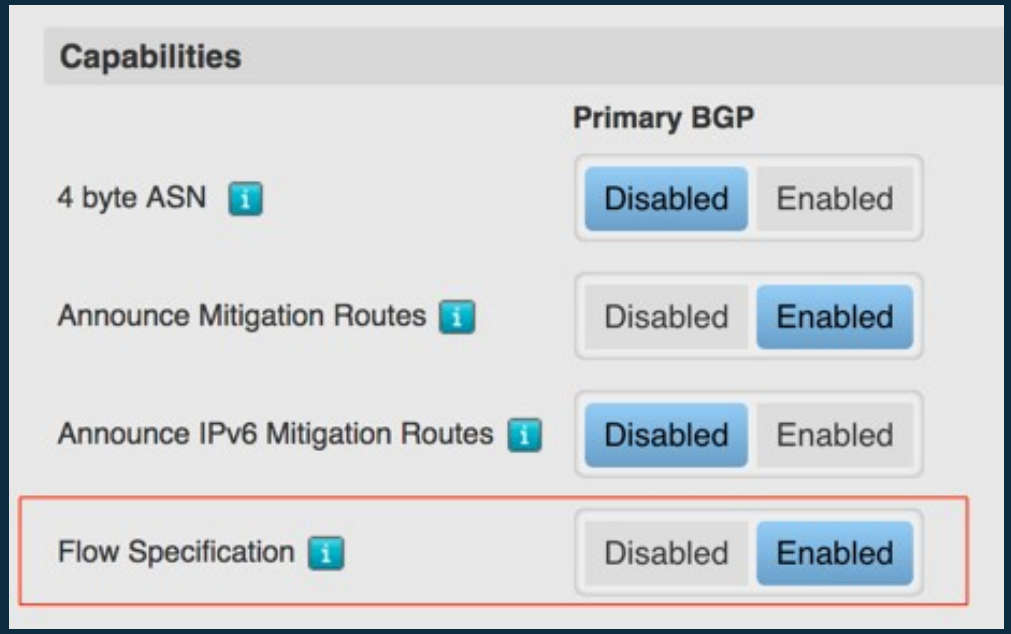

(Cisco)

```
router bgp 65555 bgp
      address-family ipv4 flowspec
        route policy FS_Policy in
        validation disable
      !
      address-family ipv6 flowspec
```

(Juniper)

```
set protocols bgp group ArborSP family inet flow no-validate Flowspec_Policy
```

Disable BGP Flowspec on an interface

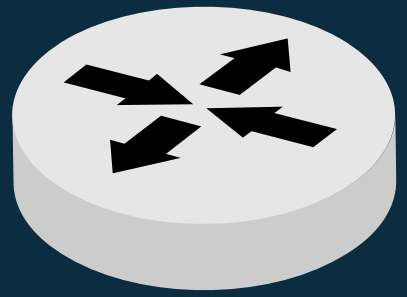

Cisco

```
flowspec local-install interface-all
!
interface TenGigE0/0/0/1
   ipv4 flowspec disable
   ipv6 flowspec disable
```

Juniper

```
set groups fs-disable interfaces ae100
set routing-options flow interface-group fs-
disable exclude
```

Validating Flowspec (Cisco)

```
RP/0/RSP0/CPU0:edge-frankfurt#show flowspec ipv4 summary
Wed Apr 12 17:44:26.947 UTC
Flowspec VRF+AFI table summary:
VRF: default
  AFI: IPv4
    Total Flows:                1
    Total Service Policies:     0
RP/0/RSP0/CPU0:edge-frankfurt#show flowspec ipv4 detail
Wed Apr 12 17:44:31.754 UTC

AFI: IPv4
  Flow            :Dest:7.7.7.7/32,Proto:=17,DPort:=0,SPort:=0
    Actions       :Traffic-rate: 0 bps  (bgp.1)
    Statistics                        (packets/bytes)
      Matched             :                   0/0
      Dropped             :                   0/0
RP/0/RSP0/CPU0:edge-frankfurt#
```

Validating Flowspec (Juniper)

```
admin@edge-tokyo> show route protocol bgp table inetflow.0

inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

7.7.7.7,*,proto=17,dstport=0,srcport=0/term:1
                    *[BGP/170] 00:00:13, localpref 100, from 172.16.1.71
                      AS path: ?, validation-state: unverified
                      Fictitious

admin@edge-tokyo> show firewall filter __flowspec_default_inet__

Filter: __flowspec_default_inet__
Counters:
Name                                                Bytes              Packets
7.7.7.7,*,proto=17,dstport=0,srcport=0                  0                    0

admin@edge-tokyo>
```

Diverting with Flowspec

Moving traffic on the network

Why change?

Flowspec Diversion Methods

IP Address

Target-VRF

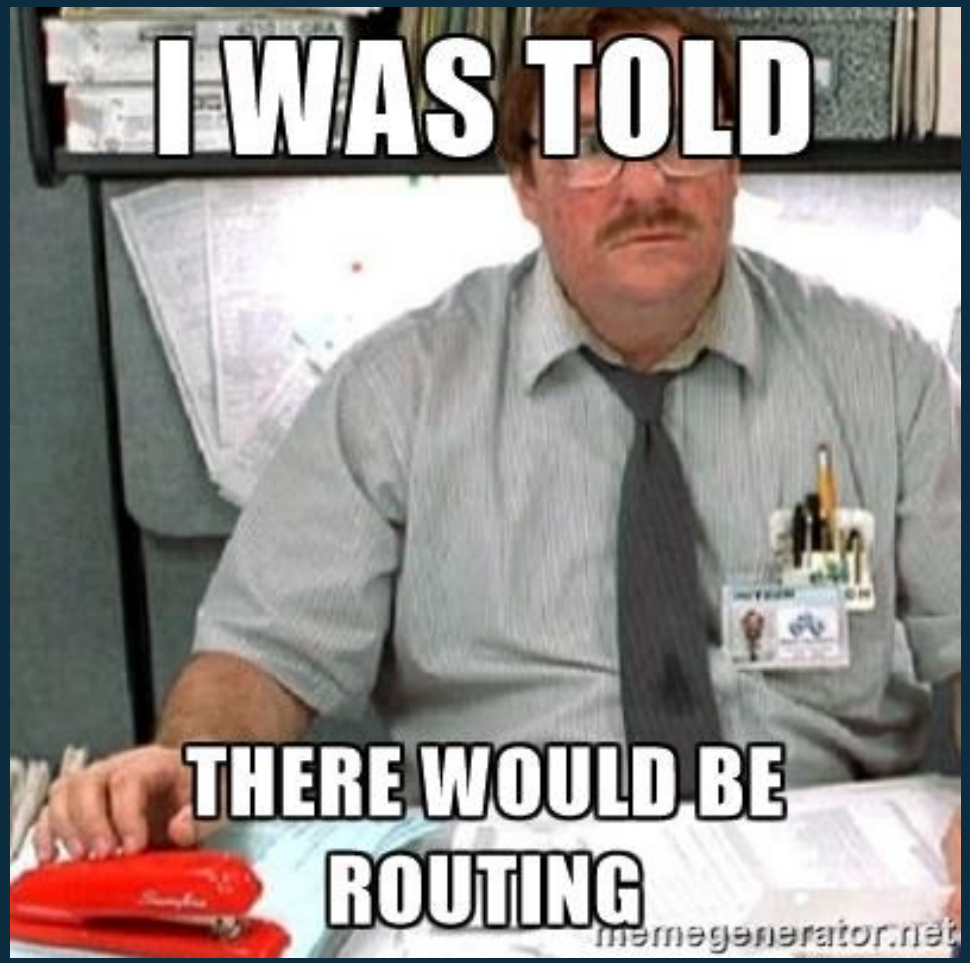

- BGP FS advertisement
- Rule: Dest-IP: 2.1.1.1 + dest-port 80
- Action: Target-vrf 65000:666
- Traffic will redirect into Dirty VRF / MPLS VPN

- Clean traffic dropped into global-routing table
- No need to encapsulate or route traffic as long as it passes non-flowspec interfaces

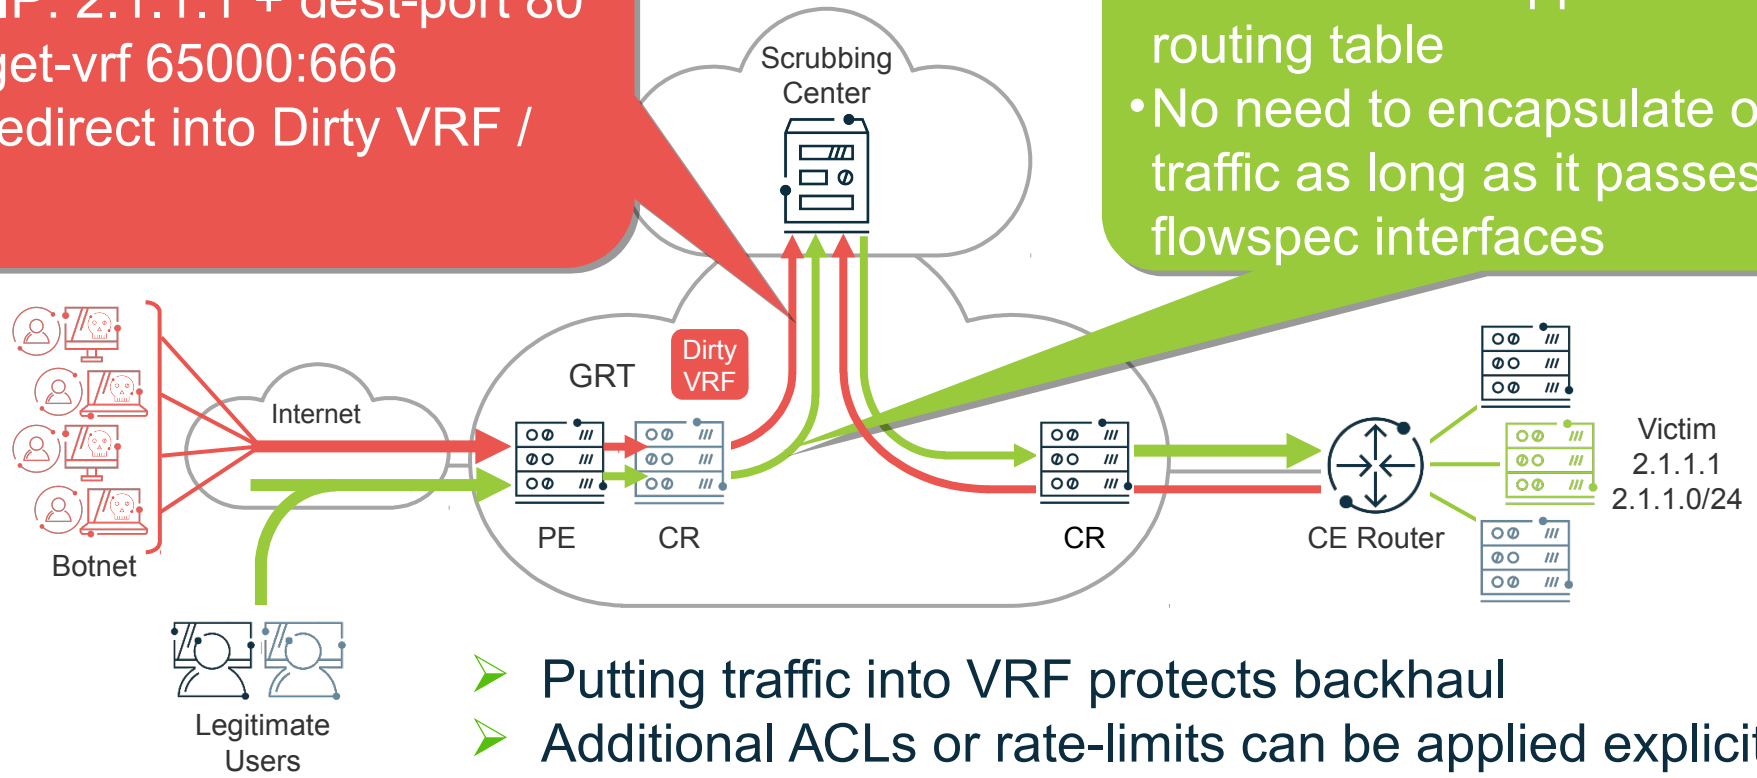

- Putting traffic into VRF protects backhaul
- Additional ACLs or rate-limits can be applied explicitly to drop likely bad traffic (SSDP, Chargen, etc.)
- Any router can put traffic into VRF/MPLS-VPN (from PE or CE side)
- Redirect can be combined with flowspec drop rules to dynamically drop known bad ports

Practice Safe Flowspec

BGP Flow specification for mitigation

Practicing Safe Flowspec

Infrastructure First

Know your victim

Understand your capabilities

Practicing Safe Flowspec

- NTP, SSDP, Chargen, MSSQL
- Careful with applications: DNS, SYN
- What is the SLA around what you are trying to protect?
  - Residential Users
  - Commercial Customers
  - Critical Infrastructure
- What services do you need to worry about?
- This is a business problem

Announcement Protection

- Respect your hardware capabilities
- Announcements
  - DDoS Mitigation Gear
  - Peers
  - Customers
  - Other use cases
- BGP policy to manage risk!

Announcement Protection

- Control rule update rates
- Prefix match validation (BGP ACLs)
- Prefix count restriction
- BGP Communities

Arbor makes it easier

- Flowspec Blacklist Offloading in TMS 8.1

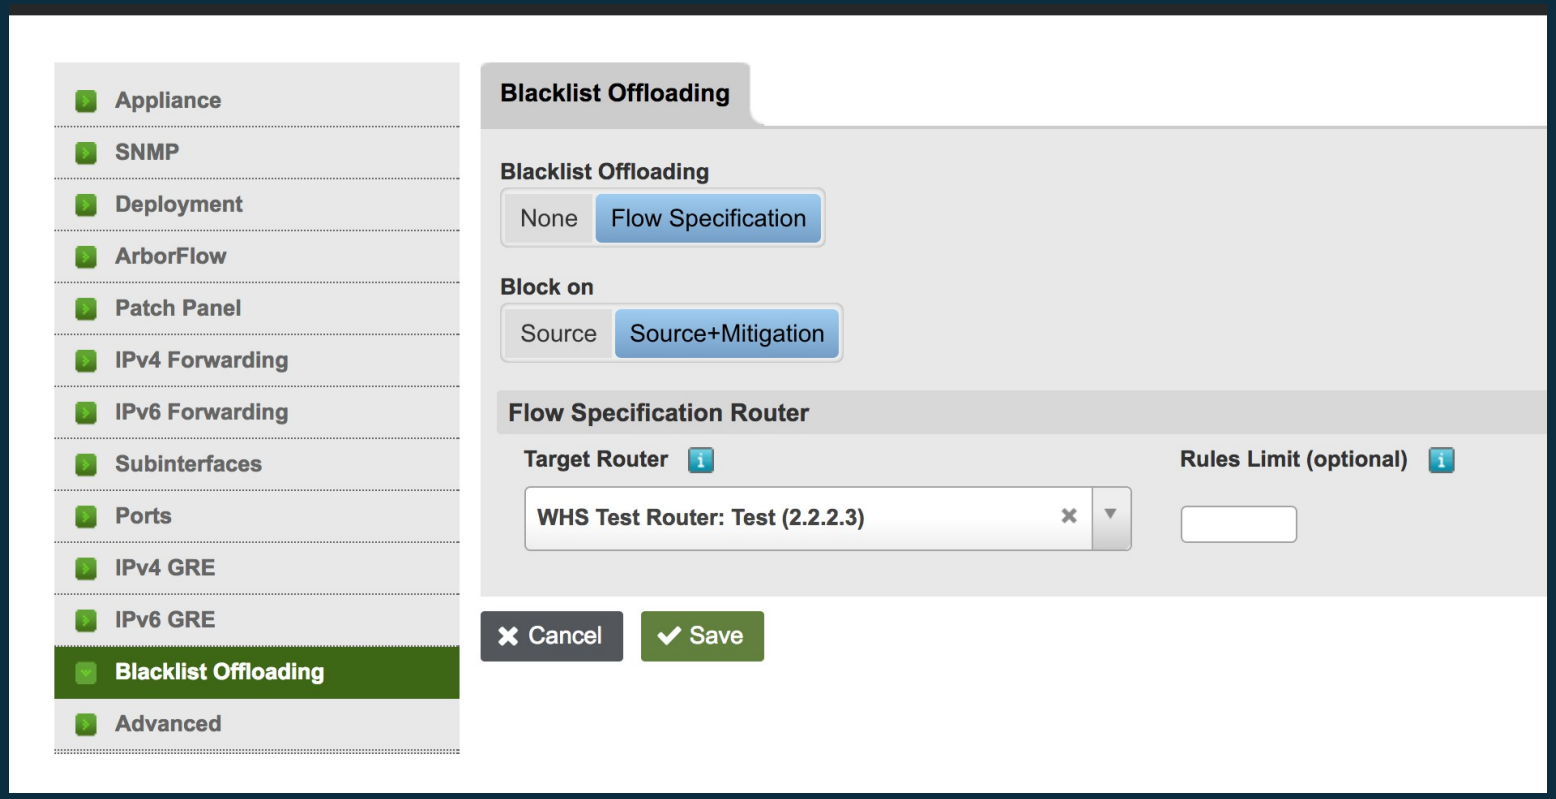

Global vs. Regional Flowspec Announcements

Announce regionally where you can!

- DDoS traffic is often not balanced
- Source addresses can vary widely
- Better scale with different announcements in different regions
- Safer when you withdraw announcements

Q&A / THANK YOU

Contact Information:

Kleber Carriello, Sr. Consulting Engineer
kco@arbor.net