NETSCOUT | Arbor

Flowspec – Just do it ?
History and Best Current Practices

LACNIC/29-FTL

Julio Arruda – Solutions Architect

Agenda

- Context
- Introduction to flowspec
- Use Case: DNS Amplification attacks mitigation
- IOT and customer x customer attacks
- Flowspec automation
- Conclusions

NETSCOUT | Arbor

Context

Background considerations

The DDoS attack surface

Any part of your network or services that is vulnerable to an attack:

- The whole PI/PA IP range
- Network interfaces
- Infrastructure
- Firewall/IPS
- Servers
- Protocols
- Applications

Attackers will find the weakness!

Traditional DDoS Mitigation

BGP Blackhole

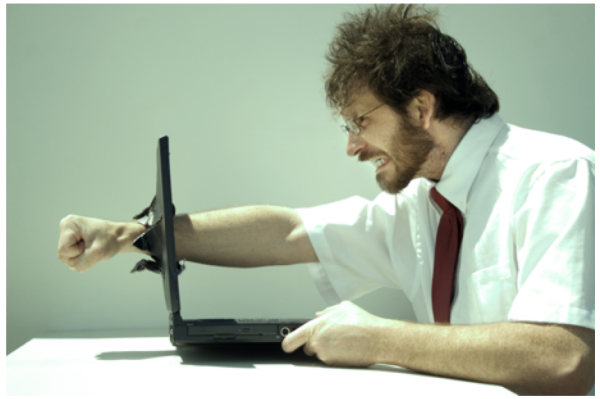

S/RTBH
Source-based /
Remotely Triggered Blackhole

IDMS
Intelligent DDoS
Mitigation System

Traditional diversion to IDMS

Based on BGP

- Legitimate and attack traffic received from Internet towards customer.
- BGP announcement is triggered pointing nexthop to a scrubbing device.
- Victim traffic is diverted to scrubbing center to identify and filter malicious traffic
- Clean traffic then gets returned to the victim via GRE, VRF or other method.

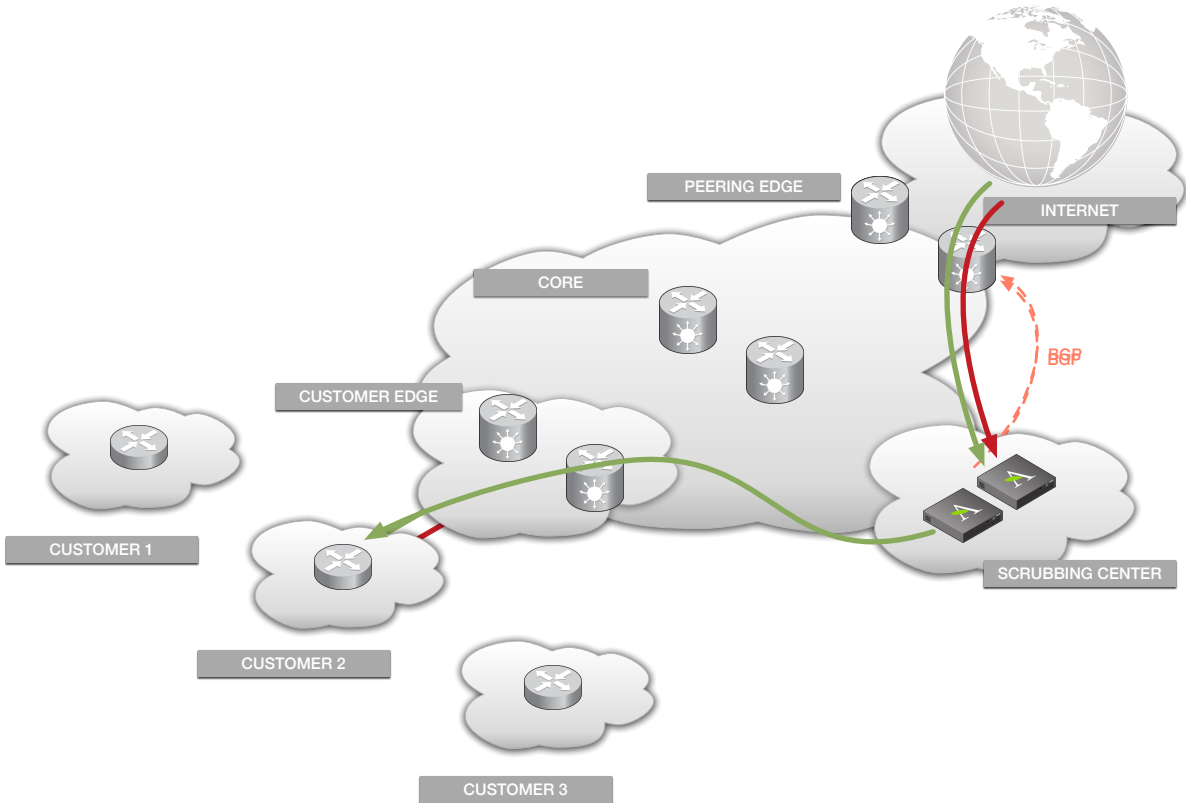

Limitations of traditional approach

There’s always room for improvement

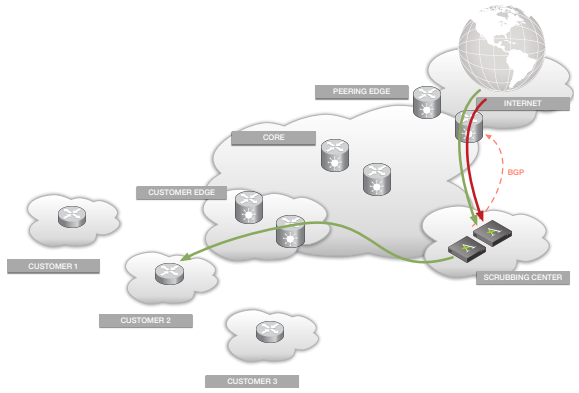

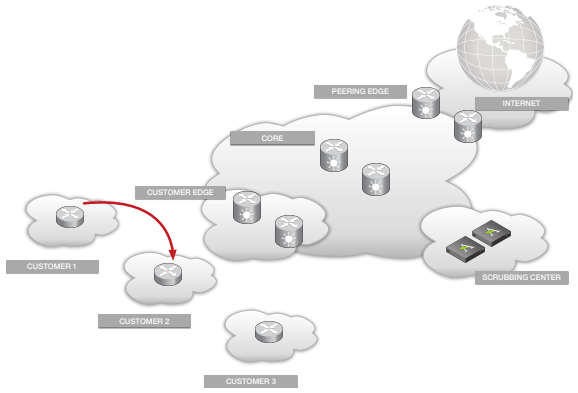

Protection Prefixes
10.245.26.14/32

?

Good for North > South attacks

Not simple for East > West attacks

All or nothing (based on hosts)

There’s a better way…

NETSCOUT | Arbor

Introduction to Flowspec

What is flowspec?

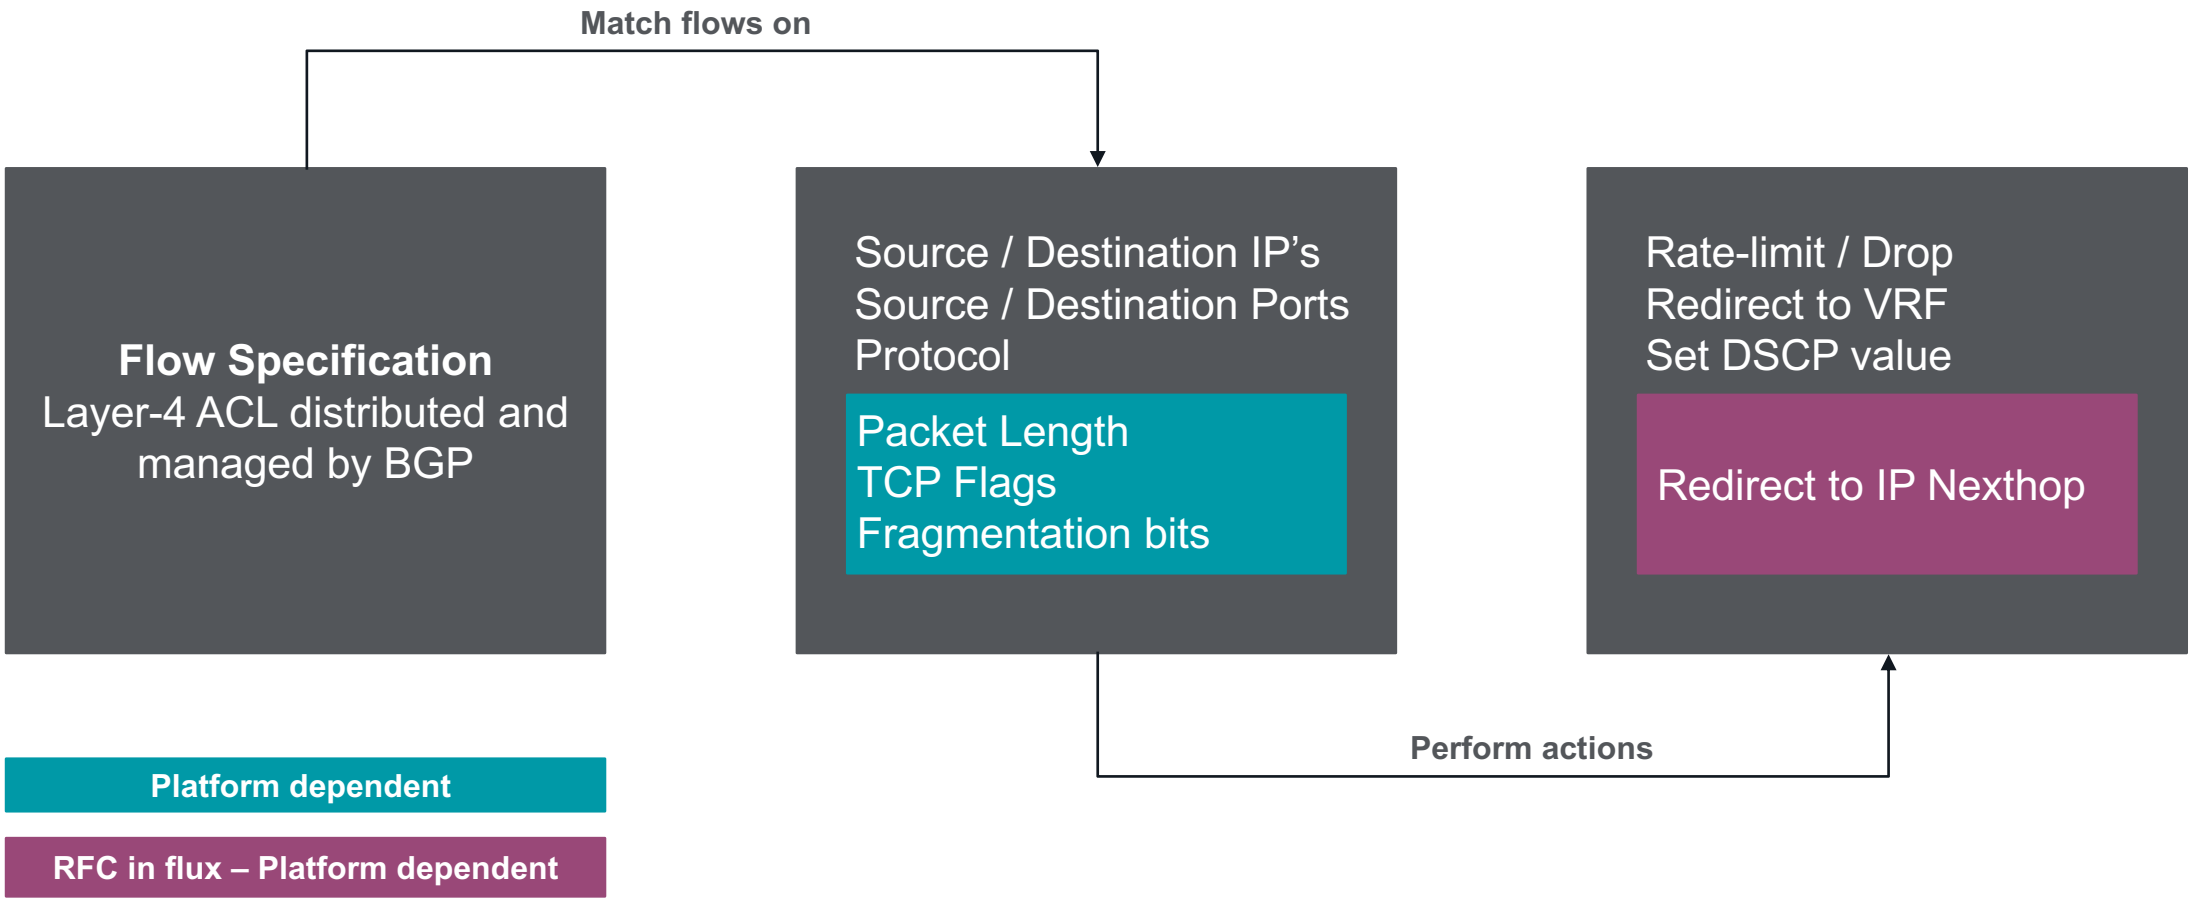

History of flowspec

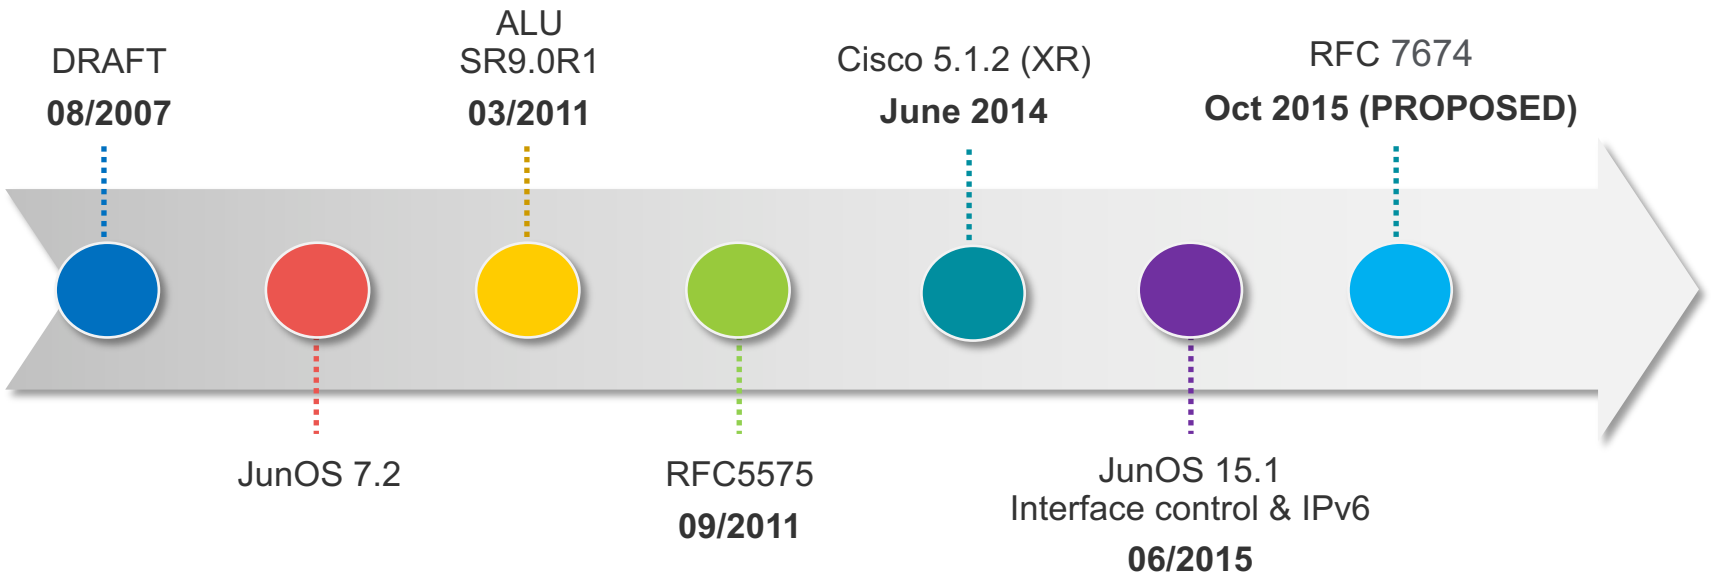

Why you should use flowspec?

Volumetric mitigation

- DDoS attack directed to a customer
- Attack traffic is detected and alerted
- Flowspec starts mitigation of attack's volumetric component
- IDMS mitigation takes care of remaining attack traffic
- IDMS delivers legitimate traffic to its destination

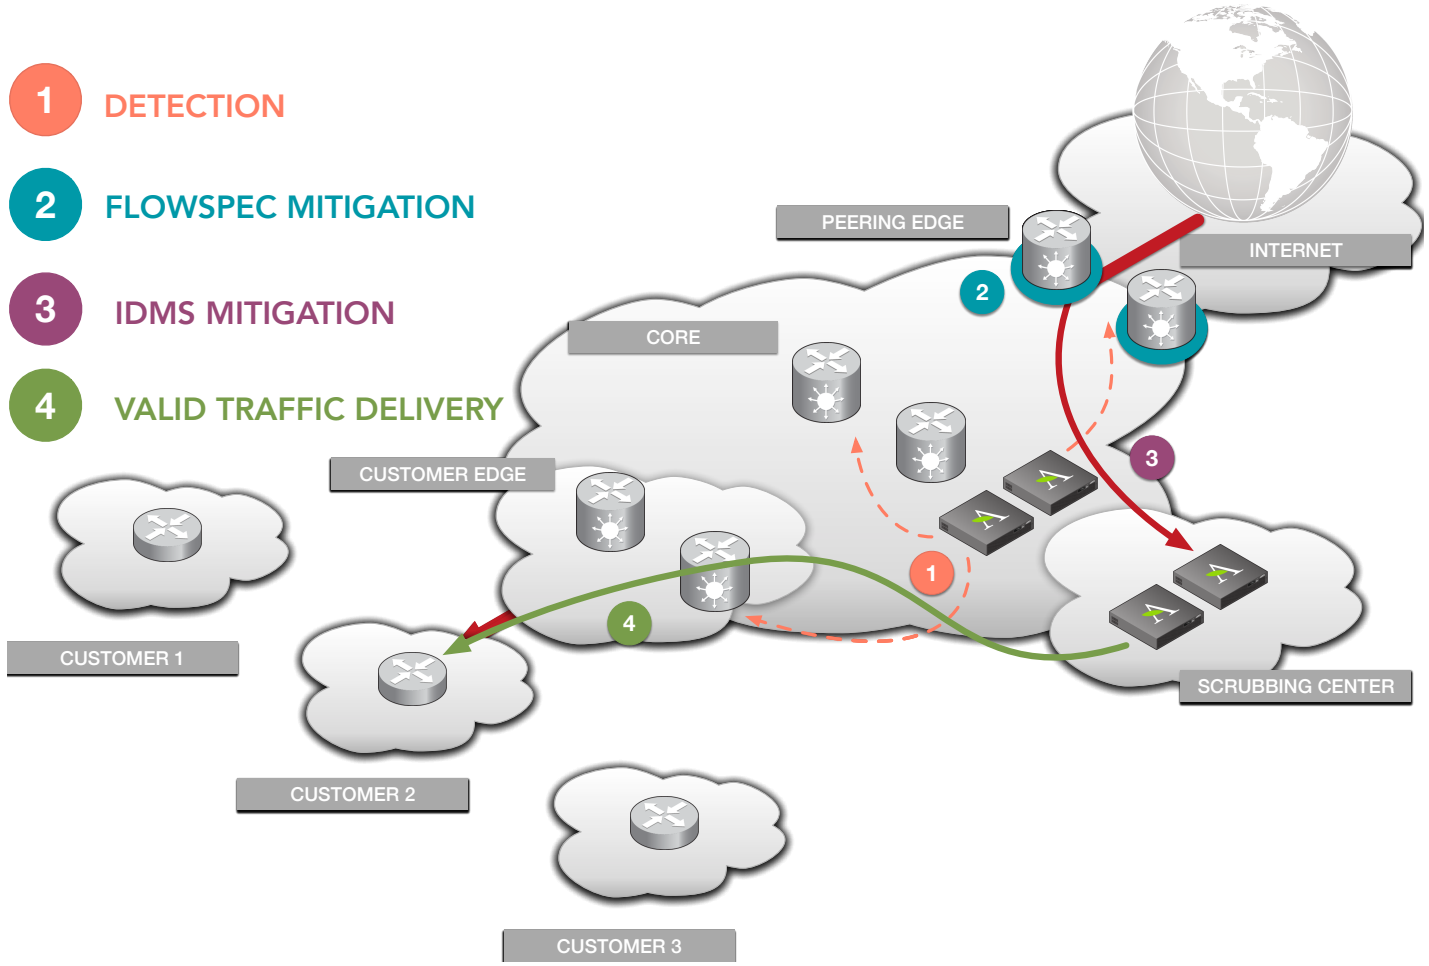

Where you should enable flowspec?

On external facing interfaces (beware of hw/sw limitations, like sub-interfaces)

Why?

- It provides ingress policy application on a router interface.
- It essentially allows PBR (Policy Based Routing)
- It specifies where flowspec rules get applied.

Benefits

- It allows flowspec rules to be applied only to untrusted places on the network (where the attack comes from).
- Removes return-traffic complexities with scrubbing centers: **No need of GRE/VRF Clean!**
- Simplifies East > West mitigation (customer to customer attacks).

NETSCOUT | Arbor

Use Case: DNS Amplification attacks mitigation

Mitigating DNS amplification – Toolset

Task: use ACLs / FlowSpec, but do not block UDP/53 completely.

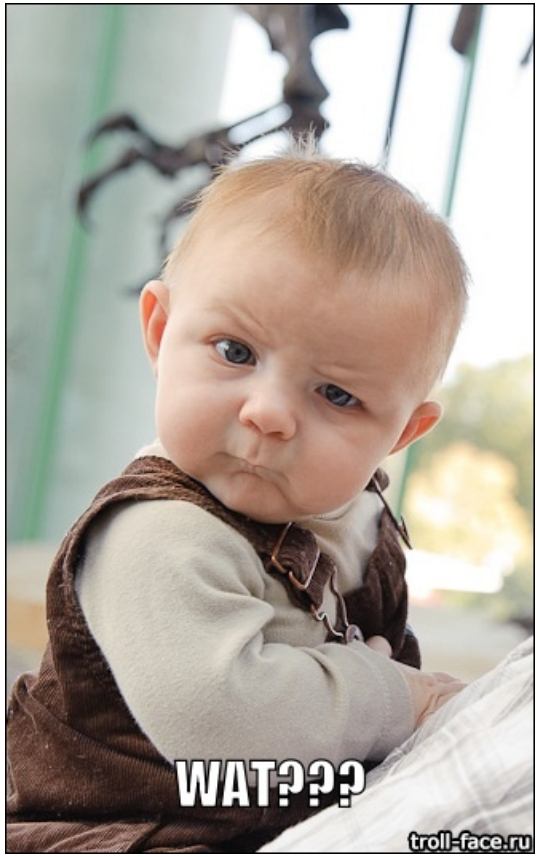

What you shouldn’t do? And why ?

Block traffic from UPD/53 completely

- It drops legitimate DNS replies
- It doesn’t drop non-initial fragments since they don’t contain UDP header
  - Amplified responses are 3-4k bytes long
  - Initial fragment is 1,500 bytes long, followed by 2-3 additional fragments
  - By blocking UDP/53 you miss **50-60%** of attack traffic

Use toolset instead!

- Run flowspec to drop initial fragments.
- Run BGP redirect to divert non-initial fragments to IDMS
- Let IDMS Invalid packets take care of non-initial fragments
- Things to keep in mind:
  - Requires IDMS capacity around **50-70%** of attack size (bps)
  - Test fragmentation bitmask before use them.

Bitmask Value	Purpose
0	Do not fragment
1	Is a fragment
2	First fragment
3	Last fragment

IOT and customer x customer attacks

Flowspec to leverage existing IDMS

Solution diagram

Enable per interface and setup “dirty” VRF

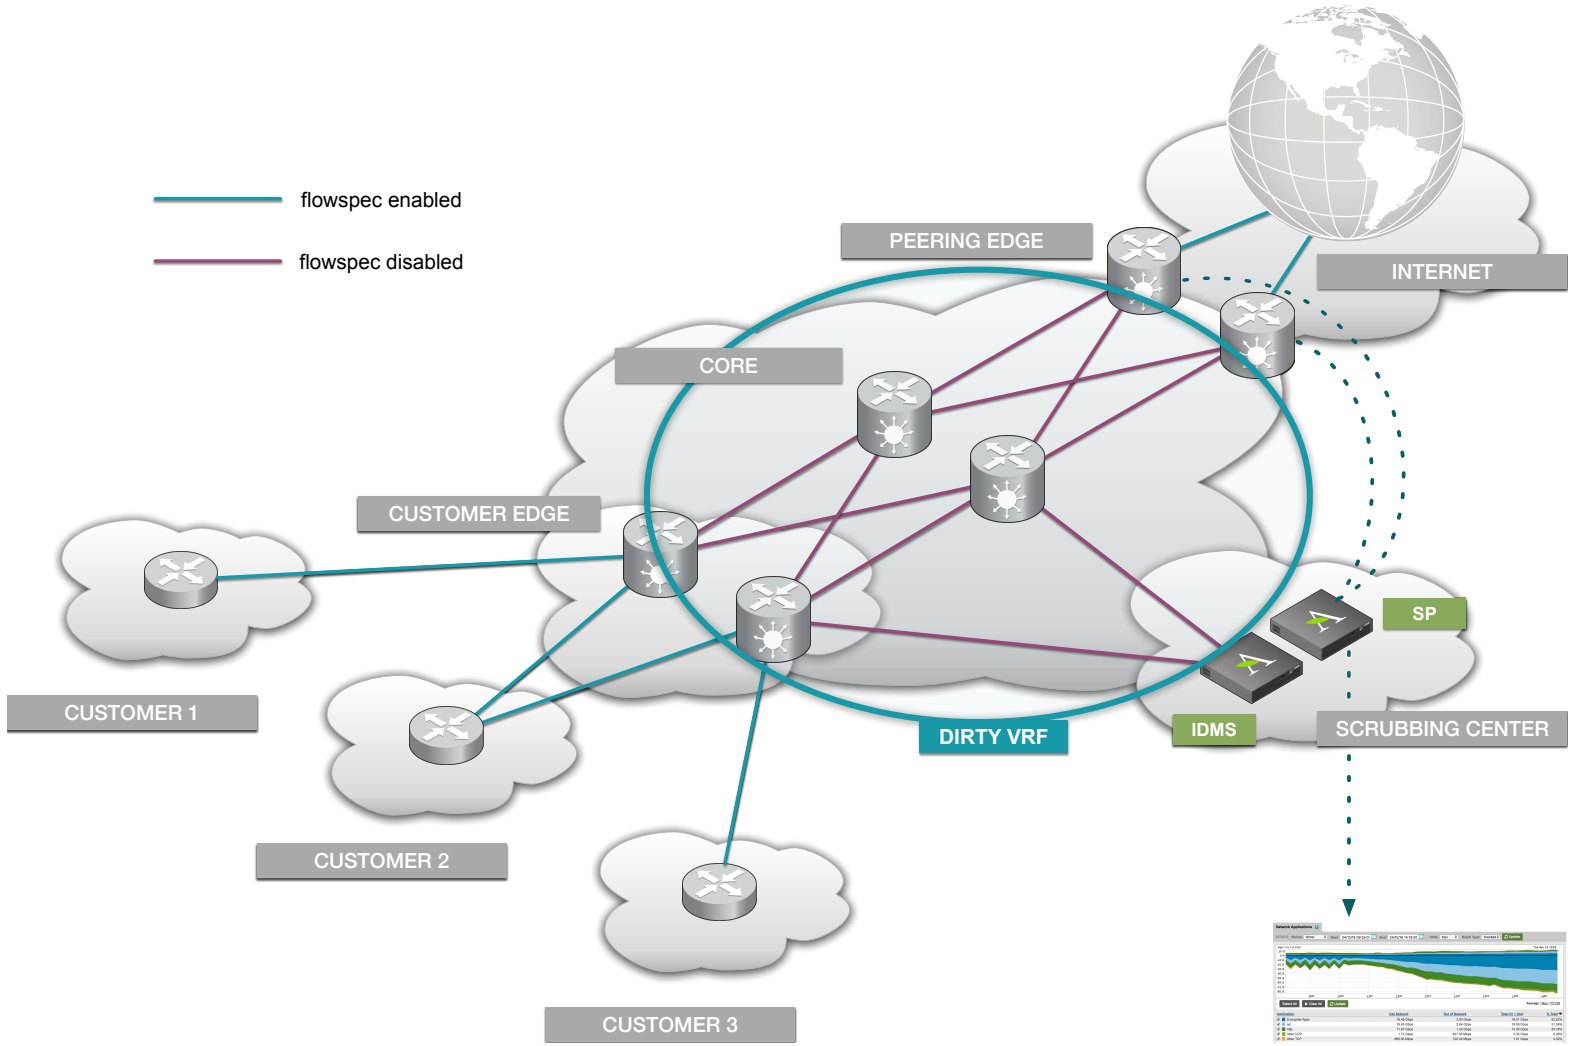

Attack scenario

North > South

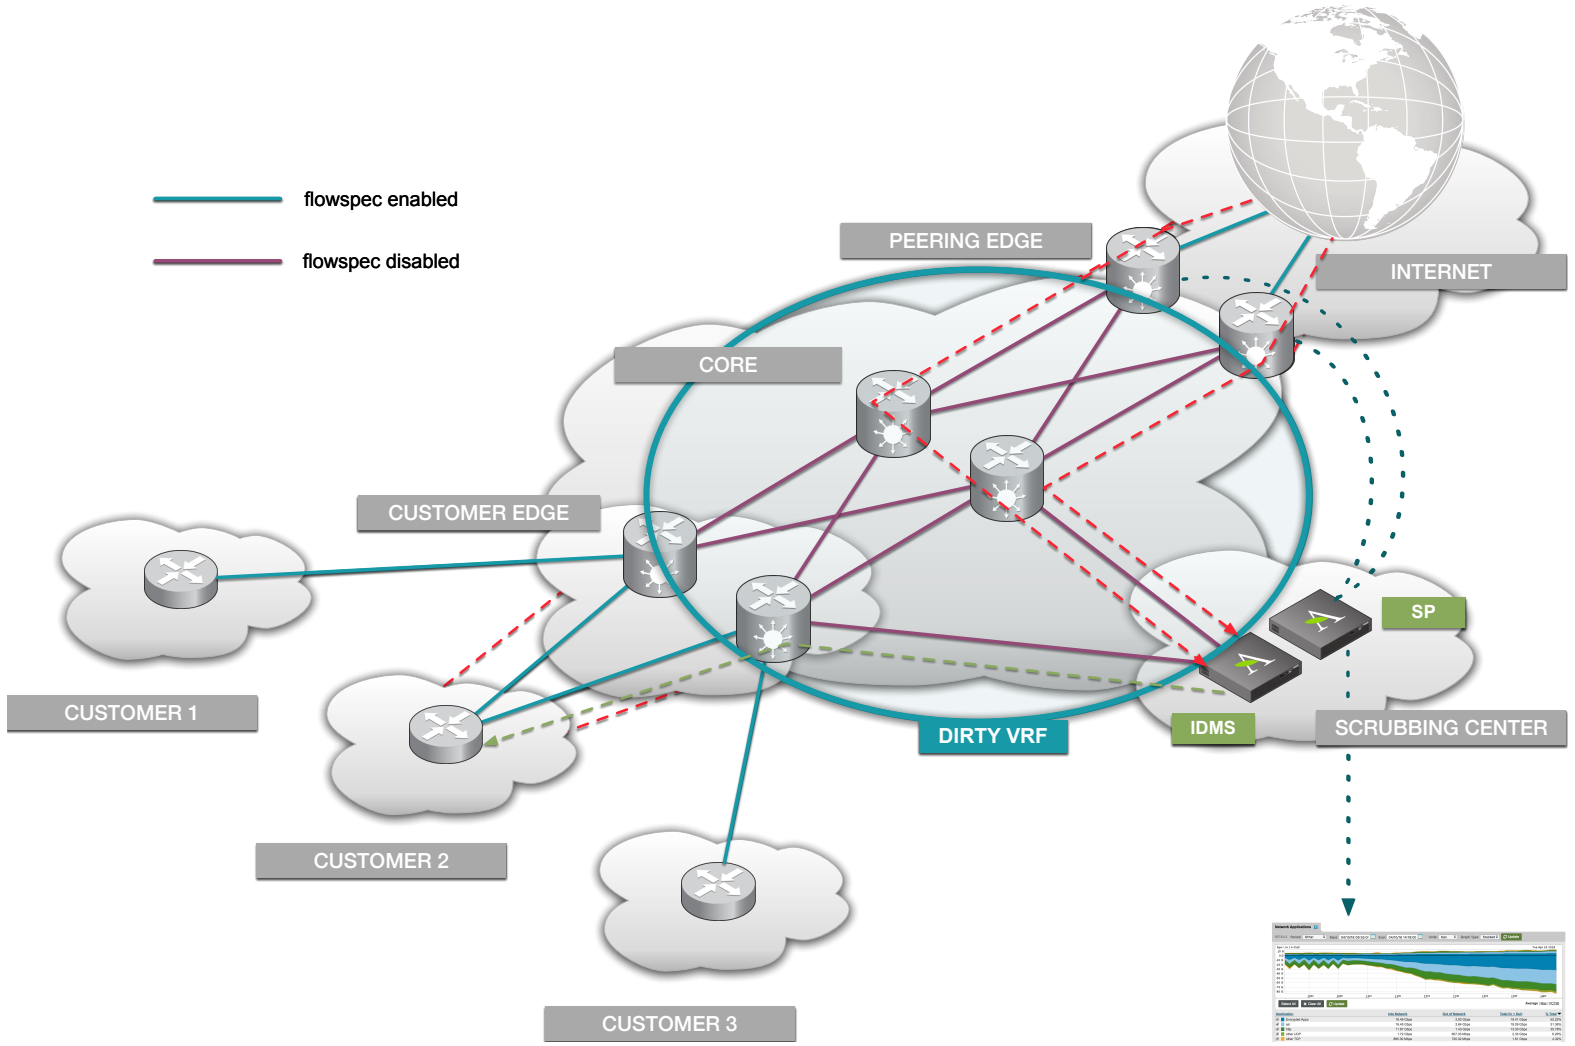

Attack scenario

East > West

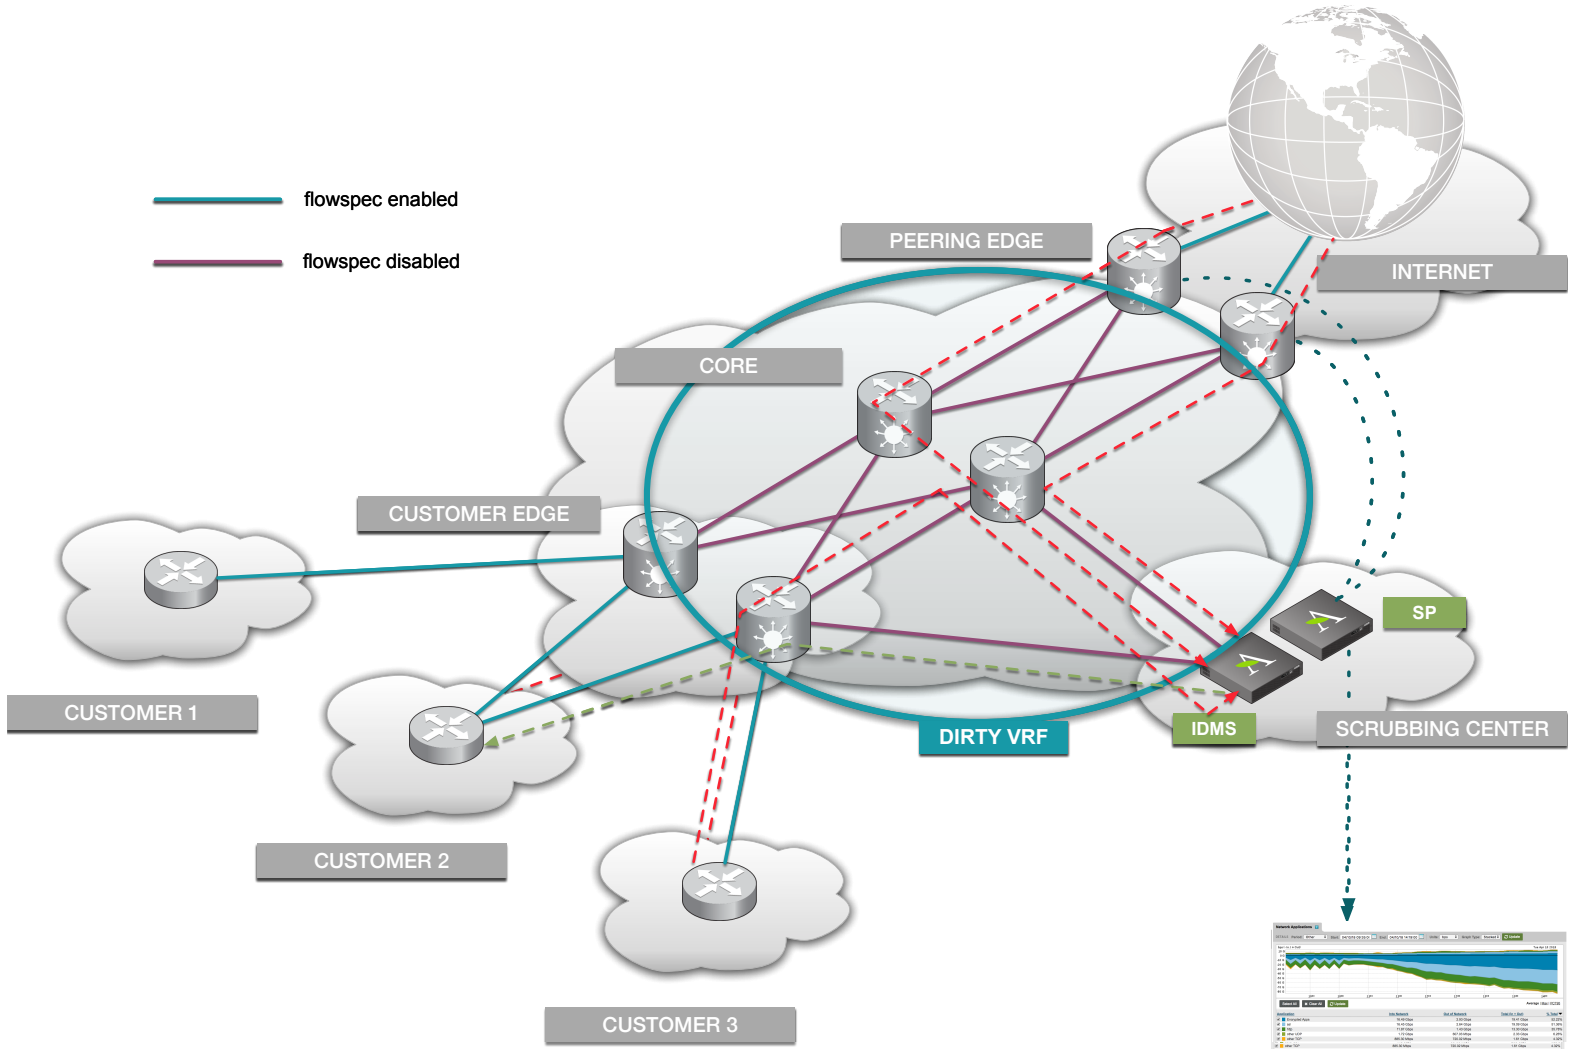

Why use ”dirty” VRF for diversion?

- VRF (Virtual Router and Forwarding) + MPLS
- Lets you further leverage network resources for attack mitigation
- Contains attack (“dirty”) traffic within a known logical entity
- Minimal routing requirements and overhead for moving traffic to scrubbers
- Easier to manage resources and protect backhaul
- With selective application of flowspec, traffic can be put back into the global routing tables without encapsulation

What if you don’t have an IDMS?

- Volumetric attack diversion is not desirable if your IDMS resources are limited or non-existent.
- You can create flowspec filters to drop both amplified responses components:
  - Initial fragments (UDP header)
  - Non-initial fragments (no UDP header)

NETSCOUT

Flowspec Automation

Previous considerations

- Safe to apply with certain protocols
- Care must be taken with others
- Think about SLA's requisites
  - Residential users
  - Enterprise customers
  - Critical infrastructure
- Identify critical services you need to be concerned about
- **Flowspec + IDMS integration adds to the solution**

Remember this is a business rather than technical problem!

Previous considerations

Continuation

- Keep in mind that flowspec typically operates at L3/L4
- Be aware that L3/L4 classification is not static (i.e. UDP/443 – QUIC)
- Ensure you clearly identify critical traffic patterns and whitelist them:
  - Name servers
  - Content Delivery Networks (CDN's)
  - Carrier-grade Network Address Translation (CGNAT)
  - Proxies

About flowspec announcements

Ensure protection when using flowspec

- Control rule update rates
- Implement prefix match validation (BGP ACL's)
- Restrict amount of announced routes
- Use BGP Communities
  - Control announcements regionally or globally
  - Tag, mark and track **who** is announcing, **what** is being announced and **where**.

If all else fails you can still use it

- Flowspec Blacklist Offloading in 8.1

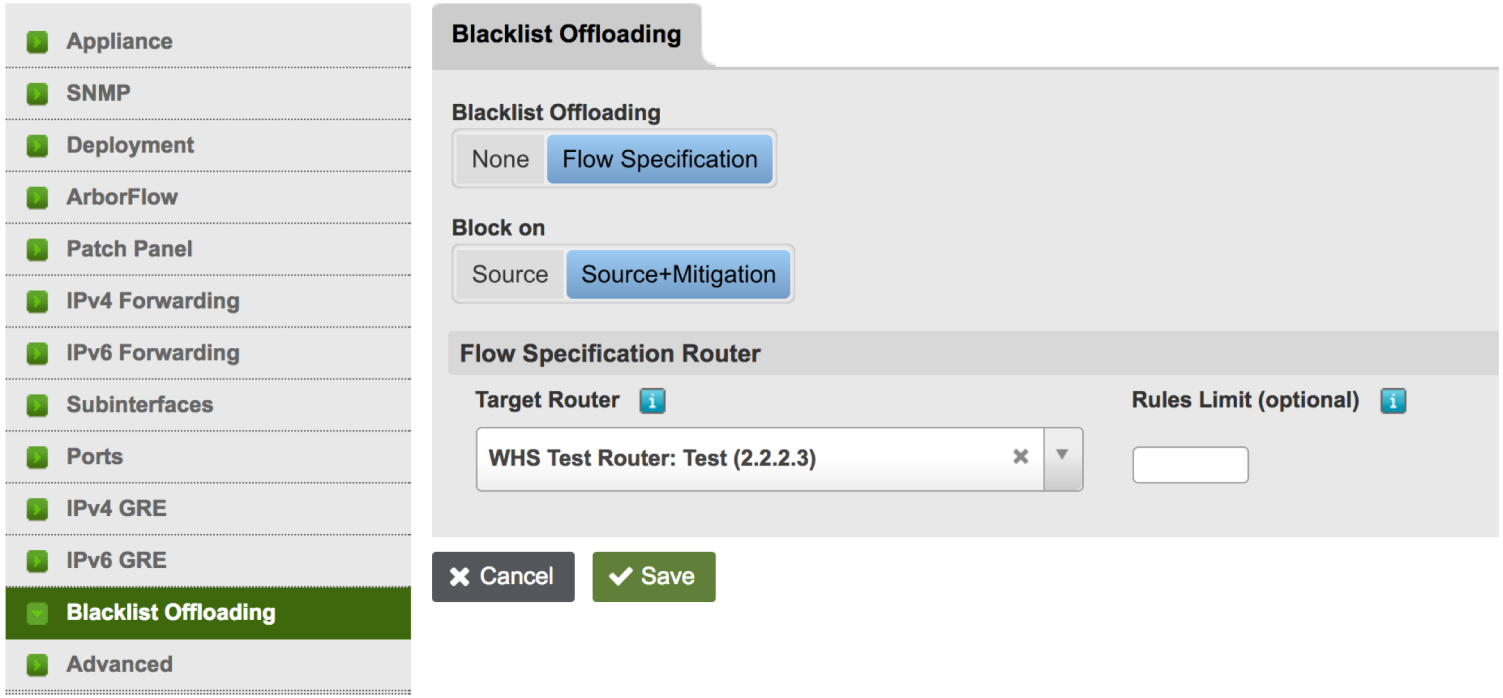

NETSCOUT

Conclusions

Just do it?

Yes, but with a clear understanding about it…

- Be aware of different functionalities between technology vendors.
- Leverage flowspec capabilities by using a single management mechanism.
- Use flowspec as part of a layered protection in conjunction with an IDMS to provide a robust security strategy.
- As ANY protocol that interfere with packet flow, establish sanity policies

Thank You.

Julio Arruda

jarruda@arbor.net

Kleber Carriello

kco@arbor.net

www.netscout.com

Example of SYN/UDP flood

Arbor Networks® SP

System | Alerts | Explore | Reports | Mitigation | Administration | Help

Thu 19 Jan 2017 19:44:10 UTC
Logged in as: tsundstrom@arbor.net (Log Out)

Security Status

Alert Activity

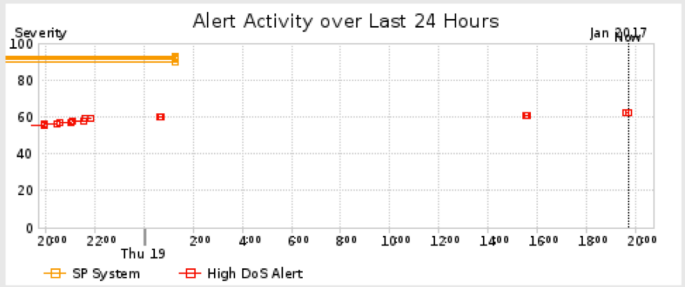

All Alerts | Alert Classes

Severity Level	Ongoing	Recent	Last 24 Hours
High	1	676	8
Medium	0	2408	4
Low	0	88	6
Total	1	3172	18

Ongoing Alerts | Ongoing Mitigations | Appliances

ID	Graph	Importance	Alert	Start Time	Classification & Annotations
13447	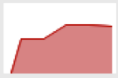	High Fast Flood 15,833.0% of 100 pps 216.3 Mbps, 28.2 Kpps	**DoS Host Alert** Incoming Host Alert to 7.7.7.7 using Victim 7 Misuse Types: IP Private, TCP SYN, Total Traffic, UDP	Jan 19 19:38 – Ongoing (0:06)	**Possible Attack** The "Total Traffic" host alert signature has been triggered at router "edge-rio". (expected rate: 100.00 Mbps/100.00 Kpps, observed rate: 198.16 Mbps/20.40 Kpps) *(by auto-annotation)*

Page generation took 0.44 seconds (Details)

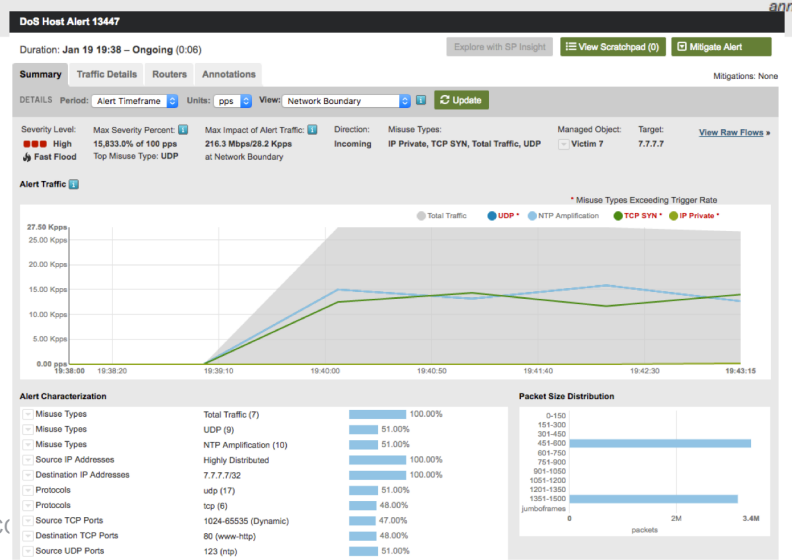

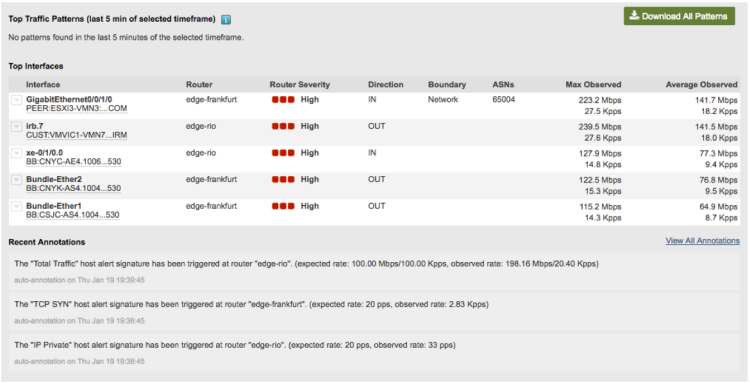

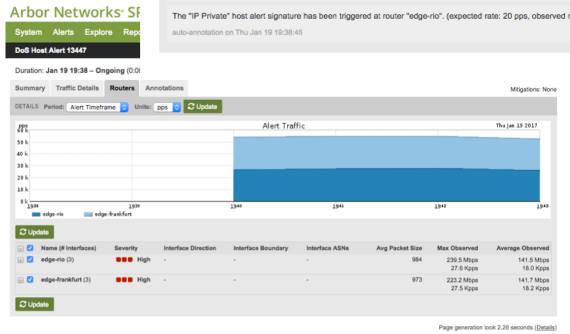

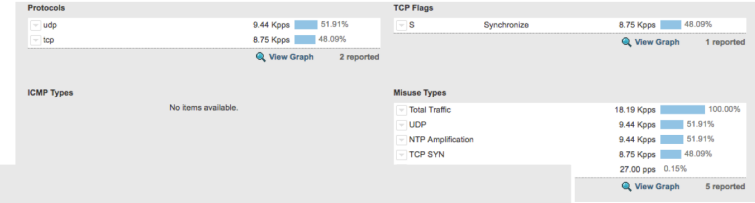

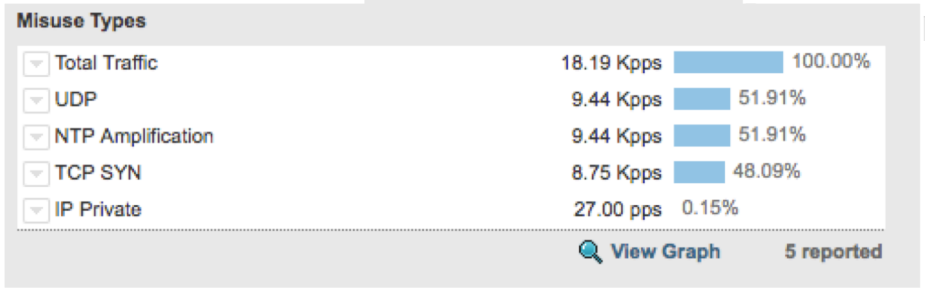

SYN, NTP amp using FS diversion and ACL

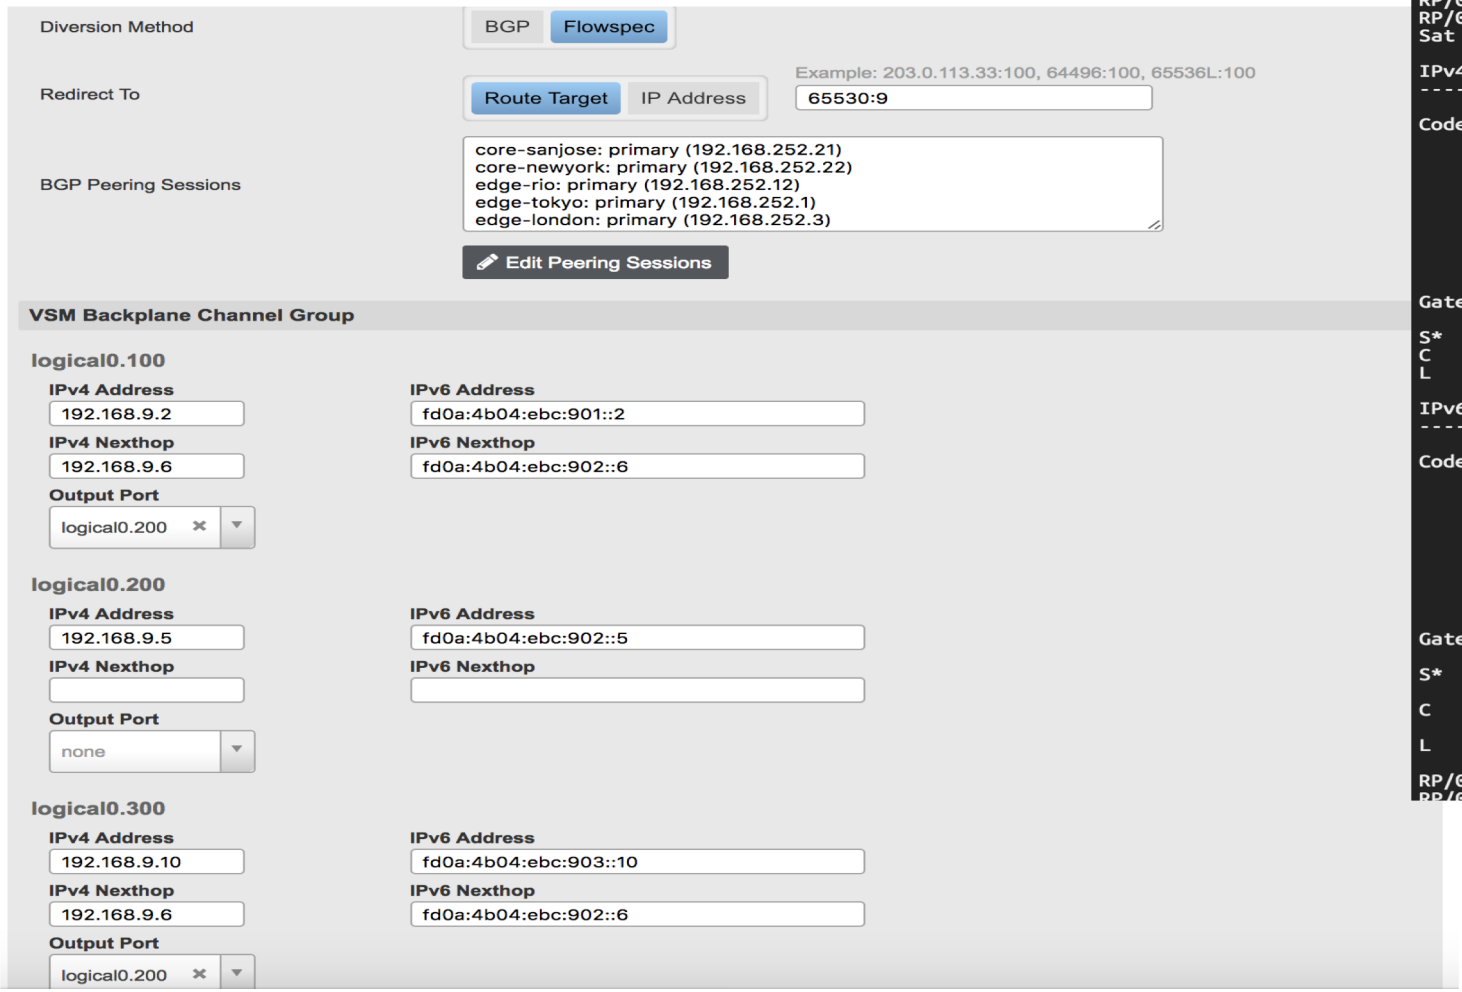


```
RP/0/RSP0/CPU0:edge-frankfurt#
RP/0/RSP0/CPU0:edge-frankfurt#
RP/0/RSP0/CPU0:edge-frankfurt#show route vrf TMS9-Dirty-VRF afi-all
Sat Jan 14 15:42:13.916 UTC

IPv4 Unicast:
-------------

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G  - DAGR, l - LISP
       A - access/subscriber, a - Application route
       M - mobile route, r - RPL, (!) - FRR Backup path

Gateway of last resort is 192.168.9.10 to network 0.0.0.0

S*   0.0.0.0/0 [1/0] via 192.168.9.10, 1w0d, Bundle-Ether9.300
C    192.168.9.8/30 is directly connected, 1w0d, Bundle-Ether9.300
L    192.168.9.9/32 is directly connected, 1w0d, Bundle-Ether9.300

IPv6 Unicast:
-------------

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G  - DAGR, l - LISP
       A - access/subscriber, a - Application route
       M - mobile route, r - RPL, (!) - FRR Backup path

Gateway of last resort is fd0a:4b04:ebc:903::10 to network ::

S*   ::/0
     [1/0] via fd0a:4b04:ebc:903::10, 1w0d, Bundle-Ether9.300
C    fd0a:4b04:ebc:903::/64 is directly
     1w0d, Bundle-Ether9.300
L    fd0a:4b04:ebc:903::9/128 is direct
     1w0d, Bundle-Ether9.300
RP/0/RSP0/CPU0:edge-frankfurt#
```

```
RP/0/RSP0/CPU0:edge-frankfurt#
RP/0/RSP0/CPU0:edge-frankfurt#
RP/0/RSP0/CPU0:edge-frankfurt#show inter bundle-ether9.300
Sat Jan 14 15:42:49.195 UTC
Bundle-Ether9.300 is up, line protocol is up
  Interface state transitions: 1
  Hardware is VLAN sub-interface(s), address is d46d.507d.6161
  Description: BB:TMS9FRA-LO0.400:AS65530 - (flowspec - redirect to RT) diversion
  Internet address is 192.168.9.9/30
  MTU 1518 bytes, BW 80000000 Kbit (Max: 80000000 Kbit)
     reliability 255/255, txload 0/255, rxload 0/255
  Encapsulation 802.1Q Virtual LAN, VLAN Id 300,  loopback not set,
  Last link flapped 1w0d
  ARP type ARPA, ARP timeout 04:00:00
  Last input 00:00:00, output 00:00:58
  Last clearing of "show interface" counters never
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 0 bits/sec, 0 packets/sec
     250660 packets input, 15308496 bytes, 0 total input drops
     0 drops for unrecognized upper-level protocol
     Received 90 broadcast packets, 0 multicast packets
     133129205 packets output, 12254258633 2 bytes, 0 total output drops
     Output 2 broadcast packets, 3789 multicast packets


RP/0/RSP0/CPU0:edge-frankfurt#
RP/0/RSP0/CPU0:edge-frankfurt#ping 192.168.9.10 vrf TMS9-Dirty-VRF
Sat Jan 14 15:42:55.583 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.9.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
RP/0/RSP0/CPU0:edge-frankfurt#
```

SYN, NTP amp using FS diversion and ACL

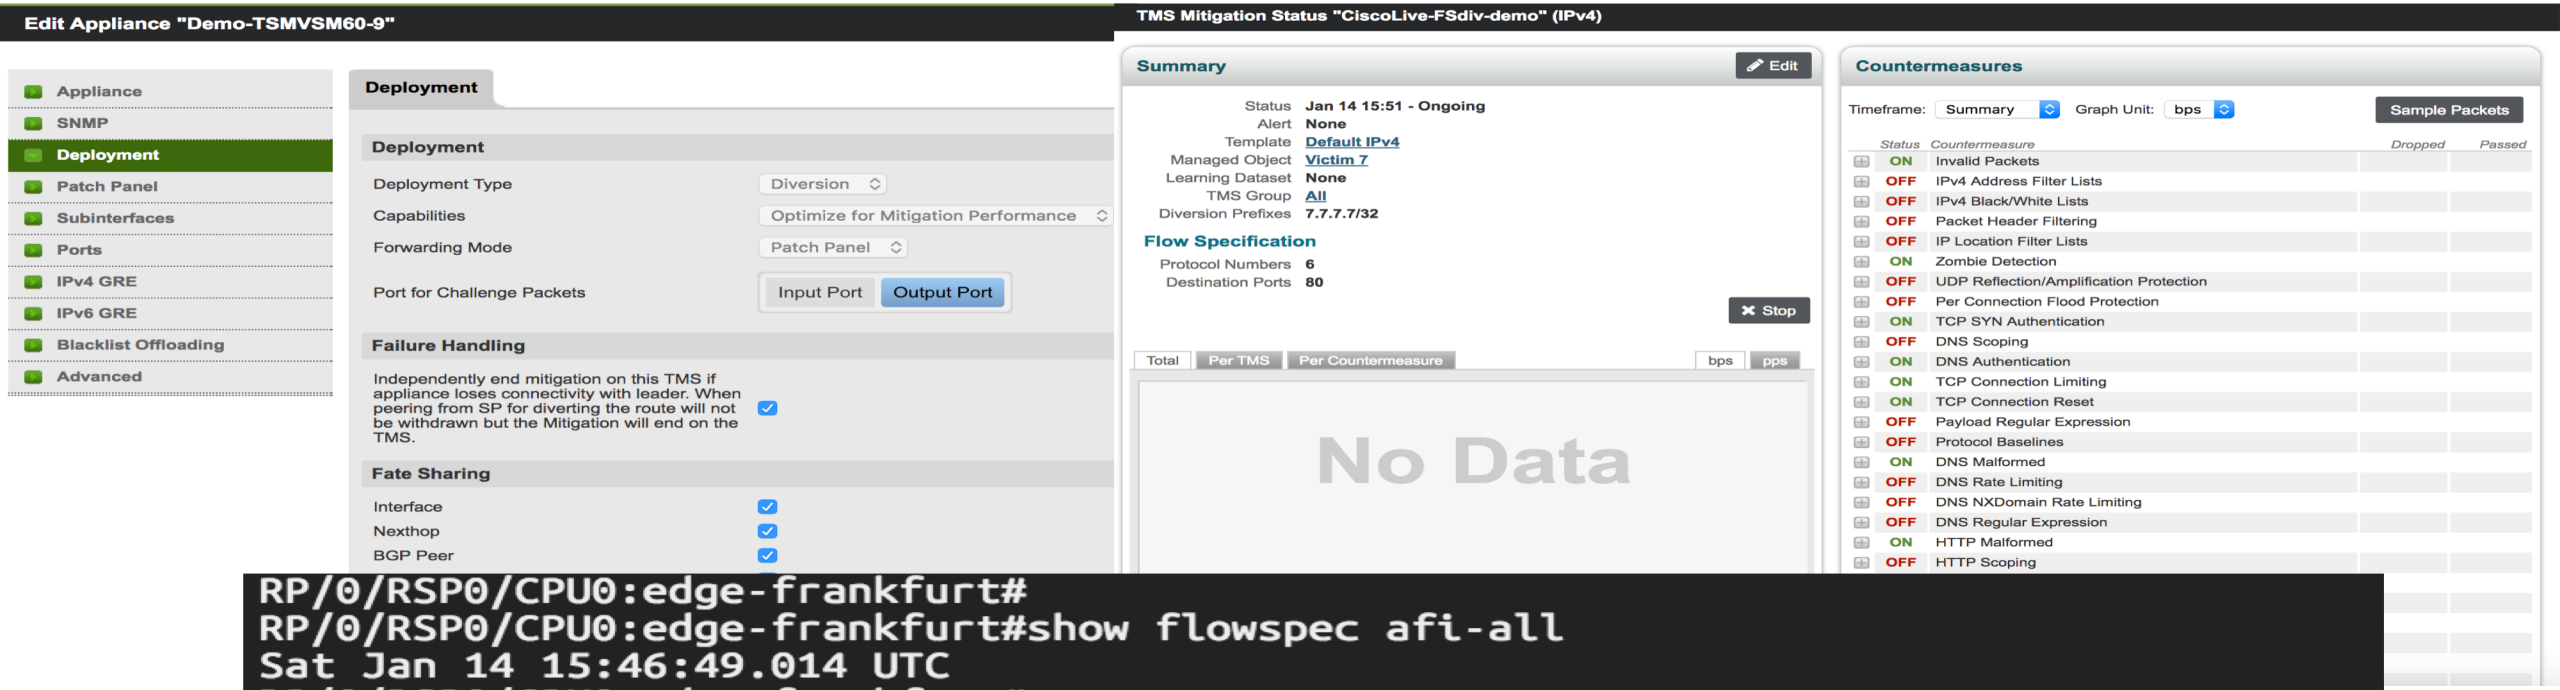


```
RP/0/RSP0/CPU0:edge-frankfurt#
RP/0/RSP0/CPU0:edge-frankfurt#show flowspec afi-all
Sat Jan 14 15:46:49.014 UTC
RP/0/RSP0/CPU0:edge-frankfurt#
RP/0/RSP0/CPU0:edge-frankfurt#show flowspec afi-all
Sat Jan 14 15:51:14.330 UTC

AFI: IPv4
  Flow                 :Dest:7.7.7.7/32,Proto:=6,DPort:=80
    Actions            :Redirect: VRF TMS9-Dirty-VRF Route-target: ASN2-65530:9  (bgp.1)
RP/0/RSP0/CPU0:edge-frankfurt#
```