IETF94 IDR WG

BGP Flowspec Interoperability Test @ Interop Tokyo 2015 ShowNet

ShowNet NOC Team member
Shuichi Ohkubo
Presenter :Cisco as ShowNet contributor

show net
Scratch and Re-build the Internet

Interop Tokyo 2015

- 8 th – 12 th June 2015
- The Number of Visitors:136,341
- Number of Exhibitors:486
- ShowNet: Interoperability test of hot topic (BGPflowspec,VXLAN/EVPN,RPKI,IEEE1588 and so on)

BGP Flowspec(RFC5575)

- Distributes ACL configuration to network routers by BGP

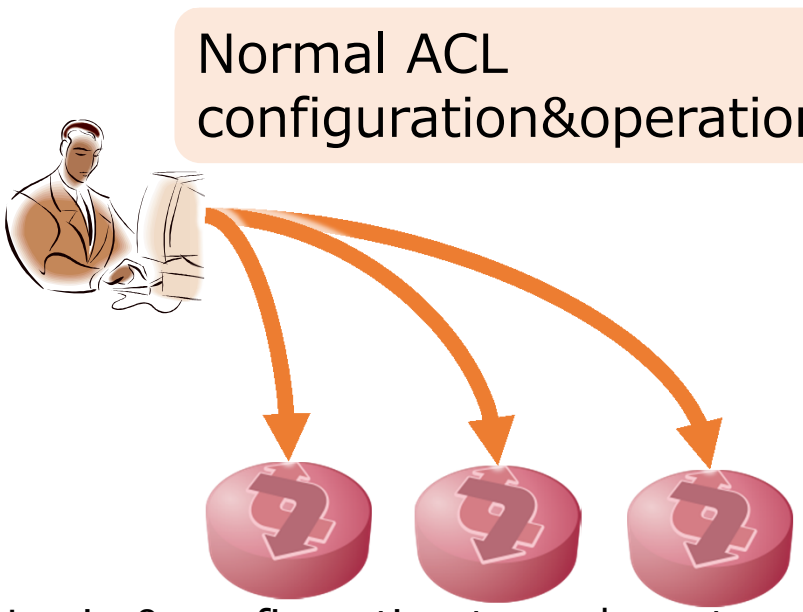

Login & configuration to each router
Too much work :(

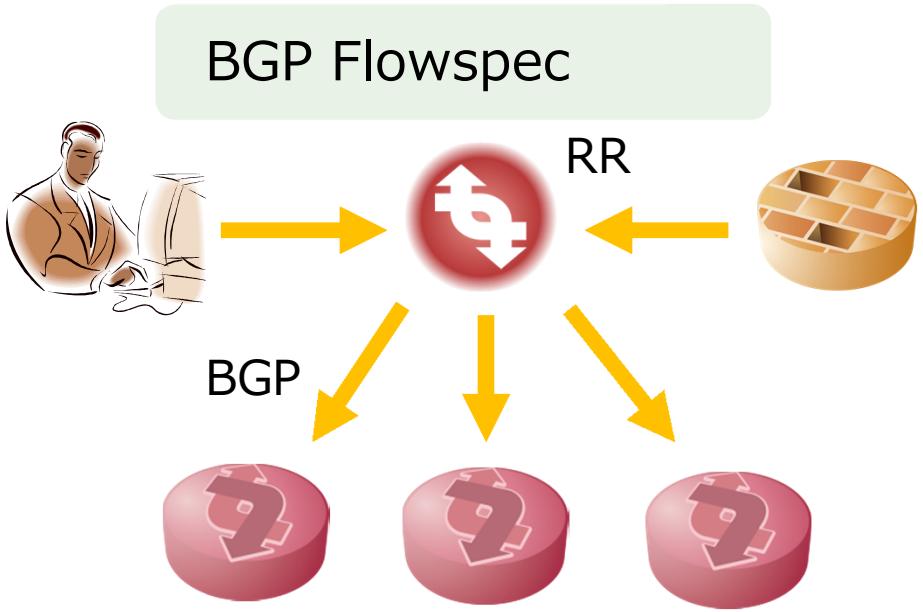

Easy to work together with security appliance

Use case

GRNET

FireCircle Operation Overview

Customer's NOC representative logs into a web tool (shibboleth) and describes flows and actions

Flow destination is validated against the customer's IP space

A dedicated router is configured (netconf) to advertise the route via BGP flowspec

eBGP sessions propagate the n-tuple to GRNET router(s). iBGP further propages the tuples to all GRNET routers.

Dynamic firewall filters are implemented on all routers

Attack is mitigated (dropped, rated-limited) upon entrance

End of attack: Removal via the tool, or auto-expire

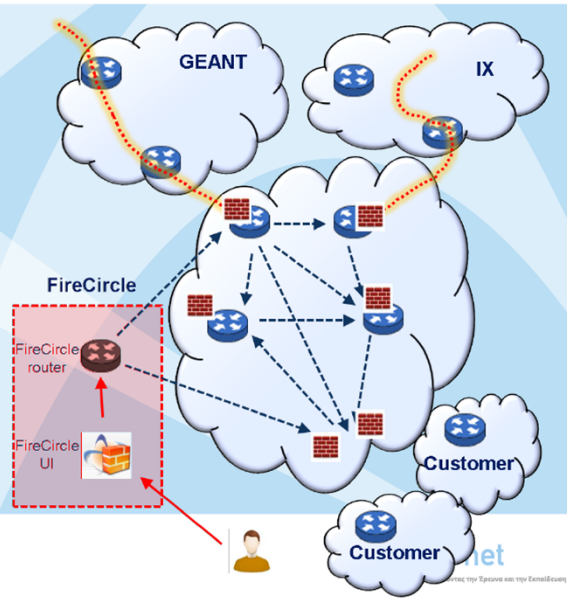

https://tnc2012.terena.org/core/presentation/41

NEO TELECOMS

Real life architecture

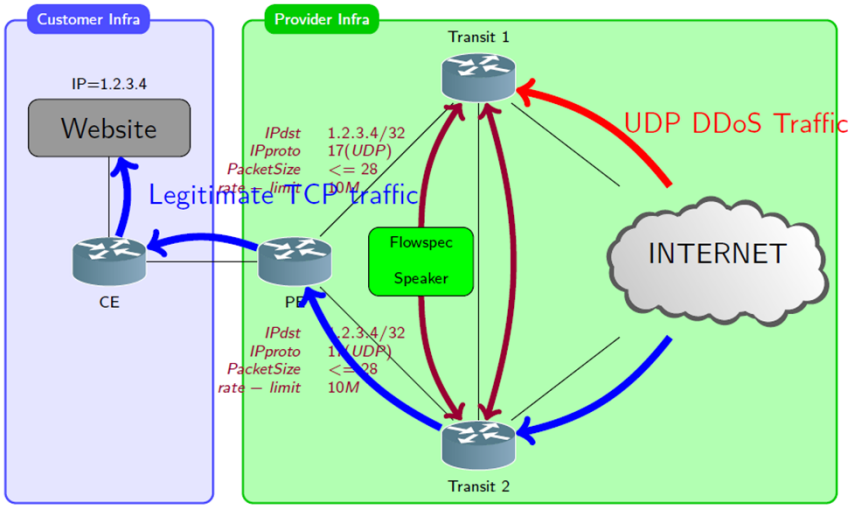

http://media.frnog.org/FRnOG_18/FRnOG_18-6.pdf

Use case

GRNET

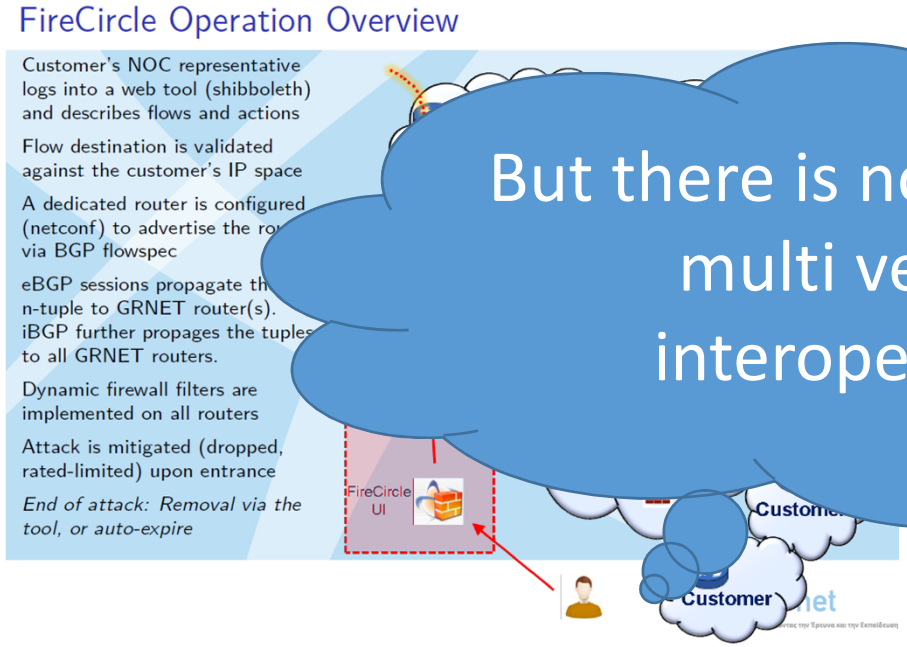

NEO TELECOM

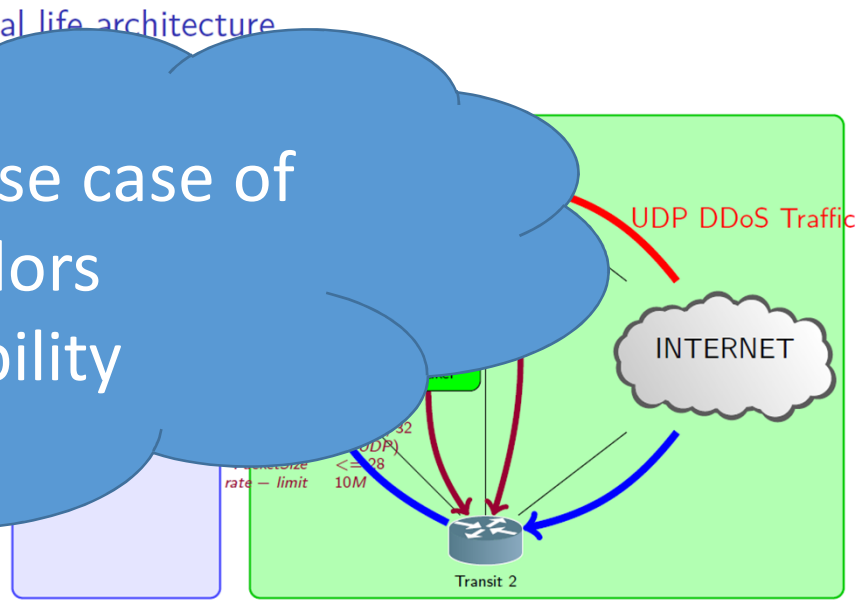

But there is no use case of multi vendors interoperability

https://tnc2012.terena.org/core/presentation/41

http://media.frnog.org/FRnOG_18/FRnOG_18-6.pdf

Interoperability test topology @Interop Tokyo 2015 ShowNet

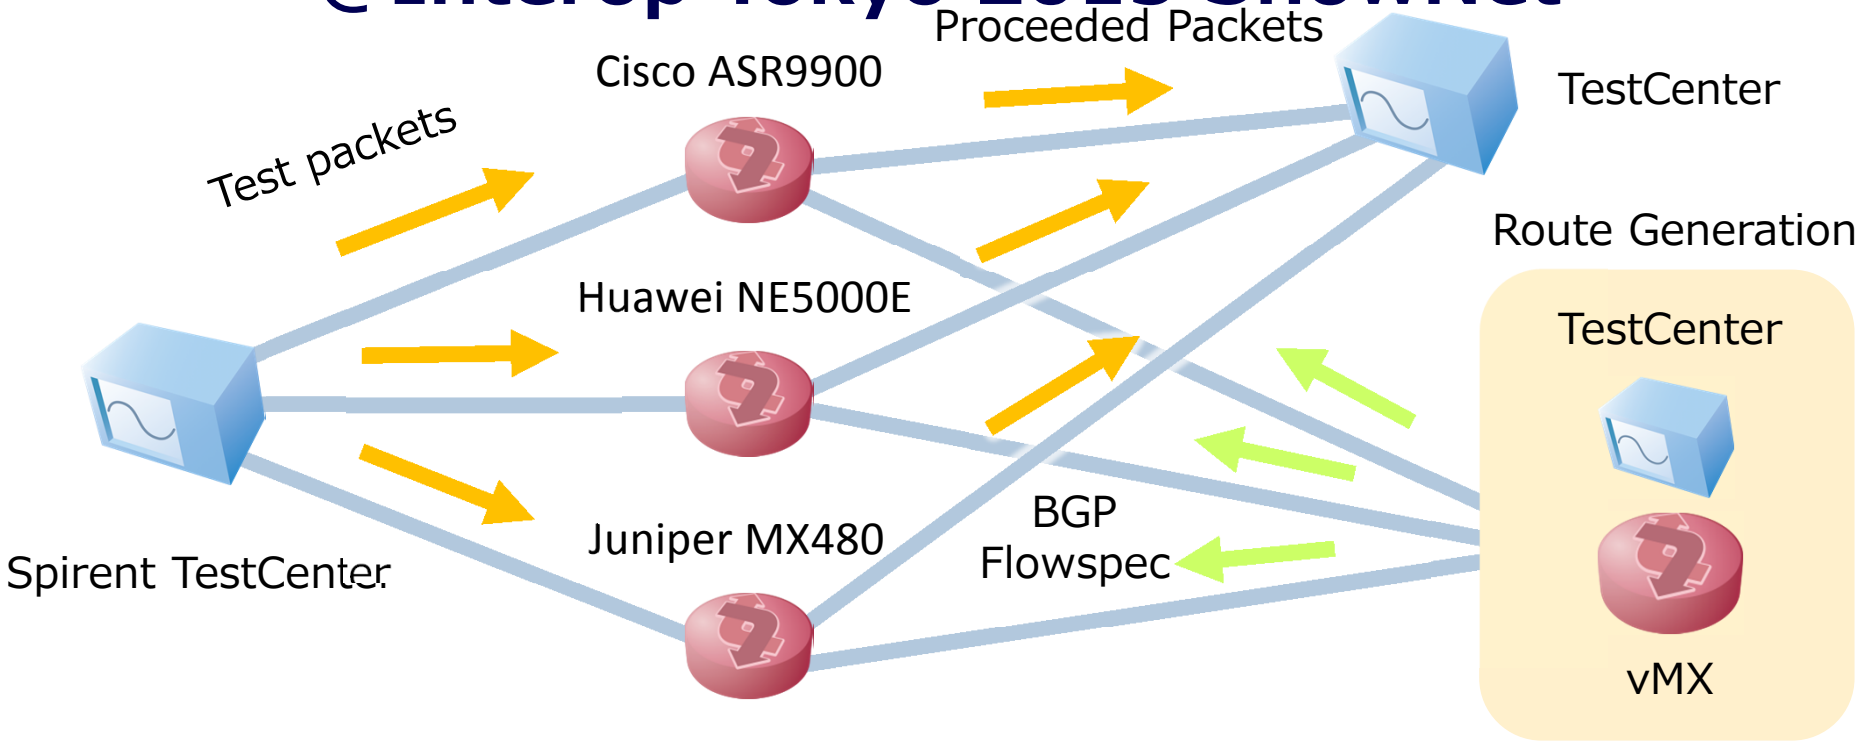

Test result BGP Flowspec Action rule

Test Item	NE5000E	ASR9900	MX480
Drop	○	○	○
Rate-limit	○	○	○
VRF Redirect	○	○	○

- Configure rate-limit=0 for Drop action
- Rate-limit: Confirmed by measuring the receiving rate to limit 100Mbps against sending 1Gbps traffic from TestCenter.
- Redirect :confirm interface counter on 3 routers and monitor latency for received packets by Spirent TestCenter

VRF Redirect

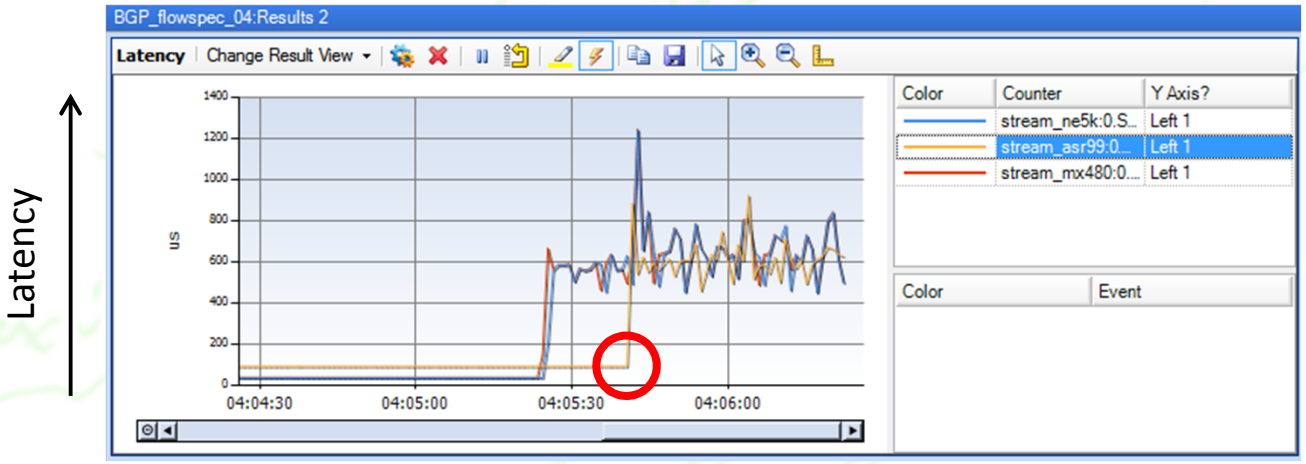

- Confirmed by measuring packets latency after redirecting (it's not caused by degradation of forwarding functionality of the router)
- ASR99xx took about 10 sec for processing after Redirection action rule injection. In case of withdrawn, the change was immediately reflected to the forwarding process.
- It depends on BGP Next-hop Scan Timer(configurable)

Rate-limit

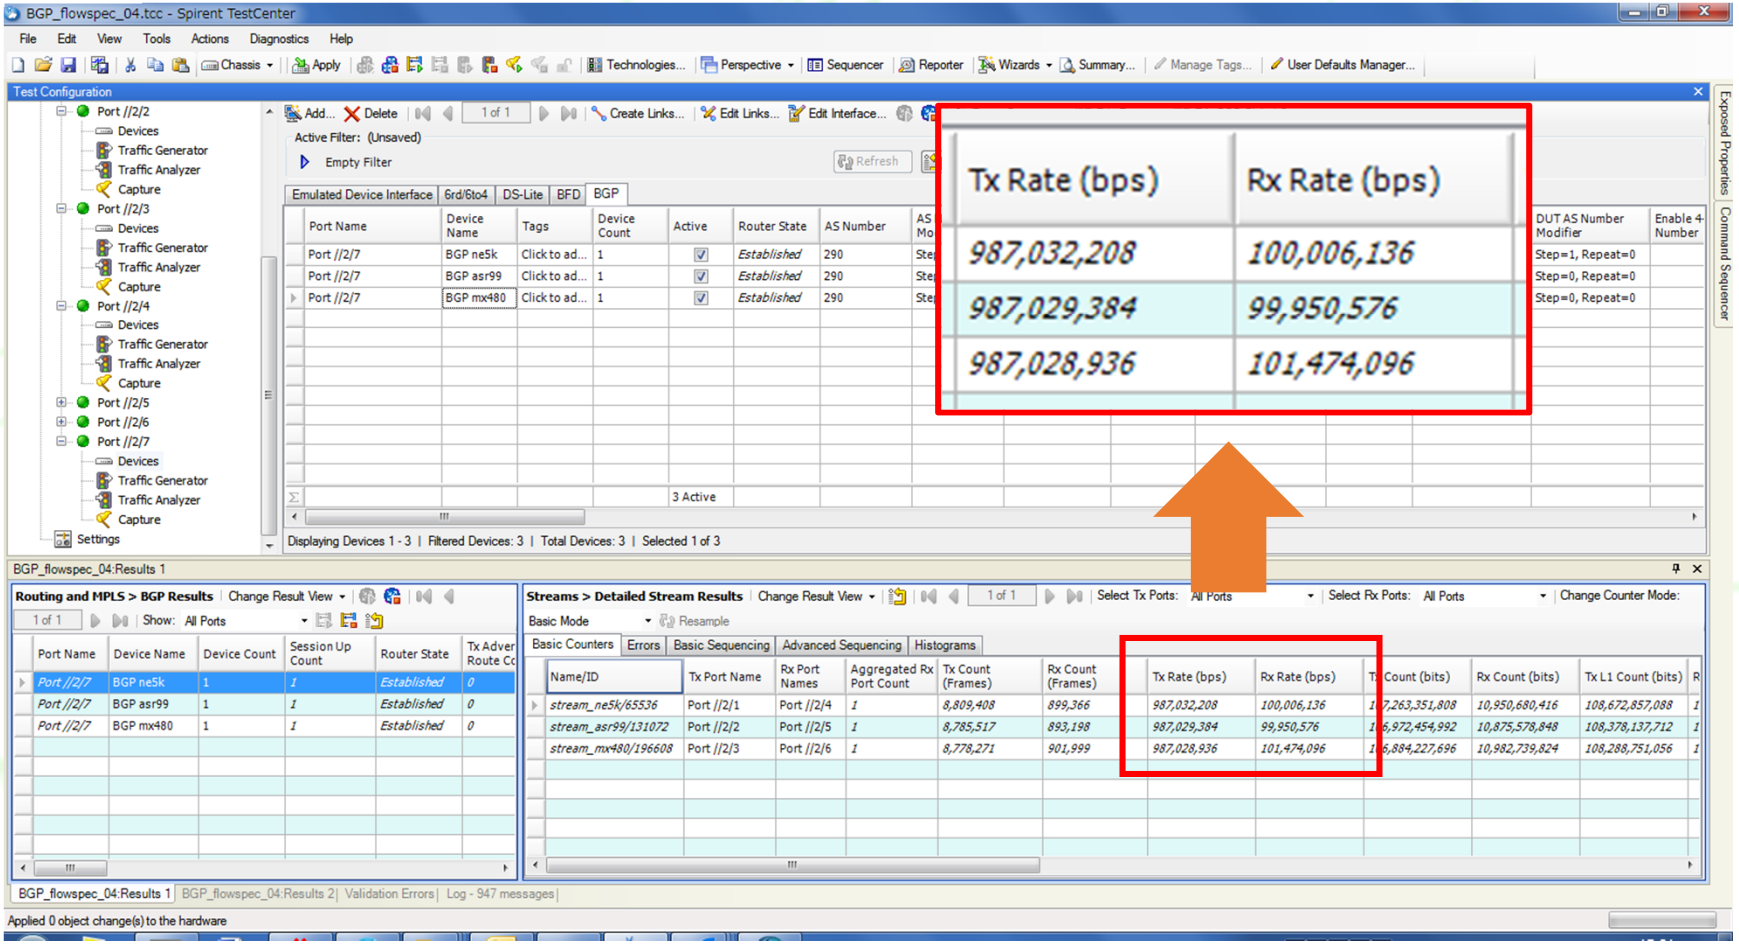
BGP_flowspec_04.tcc - Spirent TestCenter
Tx Rate (bps)
Rx Rate (bps)
987,032,208
100,006,136
987,029,384
99,950,576
987,028,936
101,474,096

Test result by Flow type

Flow type	NE5000E	ASR9900	MX480
Type 1 - Destination Prefix	○	○	○
Type 2 - Source Prefix	○	○	○
Type 3 - IP Protocol	○	○	○
Type 4 - Port	―	―	―
Type 5 - Destination port	○	○	○
Type 6 - Source port	○	○	○
Type 7 - ICMP type	○	○	○
Type 8 - ICMP code	○	○	○
Type 9 - TCP flags	○（Different NLRI）	○（Different NLRI）	○
Type 10 - Packet length	will support in Next release	○	○
Type 11 - DSCP	○	○	○
Type 12 - Fragment	―（Different NLRI）	○	○

Difference in NLRI format Type9. TCP Flags

Juniper

Configure syn+ack

Dest	/32	45.0.2.54	Src	/32	45.0.2.42	TCP Flg.	op	Bit mask	op	Bit mask
0x01	20	2d000236	02	20	2d00022a	09	00	02	80	10

0x02 SYN

0x10 ACK

Cisco

Dest	/32	45.0.2.54	Src	/32	45.0.2.42	TCP Flg.	op	Bit mask
0x01	20	2d000236	02	20	2d00022a	09	81	12

0x12
ACK-SYN

Difference NLRI format
Type9. TCP Flags

ASR receives NLRI but does not work as expected

Cisco provides special firmware during the Interop period
, confirmed work as expected
(It's already integrated in 5.3.2 as CSCuu79956)

Difference in Match bit Type9. Type12.

Juniper

op=0x80

```
  0   1   2   3   4   5   6   7
+---+---+---+---+---+---+---+---+
| e | a |  len  | 0 | 0 |not| m |
+---+---+---+---+---+---+---+---+
  1   0   0   0   0   0   0   0
```

Cisco, Huawei

op=0x81

```
  0   1   2   3   4   5   6   7
+---+---+---+---+---+---+---+---+
| e | a |  len  | 0 | 0 |not| m |
+---+---+---+---+---+---+---+---+
  1   0   0   0   0   0   0   1
```

NE5000E treat as Invalid m=0

Huawei will support in future (support m=0)

Operation example on ShowNet

Always seen SSH Brute-force attack
to ShowNet

Execute filtering by BGP Flowspec

1. permit TCP Port 22 from specific server
2. drop 45.0.0.0/16 TCP Port 22,23

order of evaluation is important

Need additional command for JUNOS

```
set routing-options flow term-order standard
```

http://www.juniper.net/documentation/en_US/junos14.2/topics/topic-map/bgp-flow-routes.html

By default, the Junos OS uses the term-ordering algorithm defined in version 6 of the BGP flow specification draft. In Junos OS Release 10.0 and later, you can configure the router to comply with the term-ordering algorithm first defined in version 7 of the BGP flow specification and supported through RFC 5575, Dissemination of Flow Specification Routes.

Best Practice: We recommend that you configure the Junos OS to use the term-ordering algorithm first defined in version 7 of the BGP flow specification draft. We also recommend that you configure the Junos OS to use the same term-ordering algorithm on all routing instances configured on a router.

Combination demo with SAMURAI

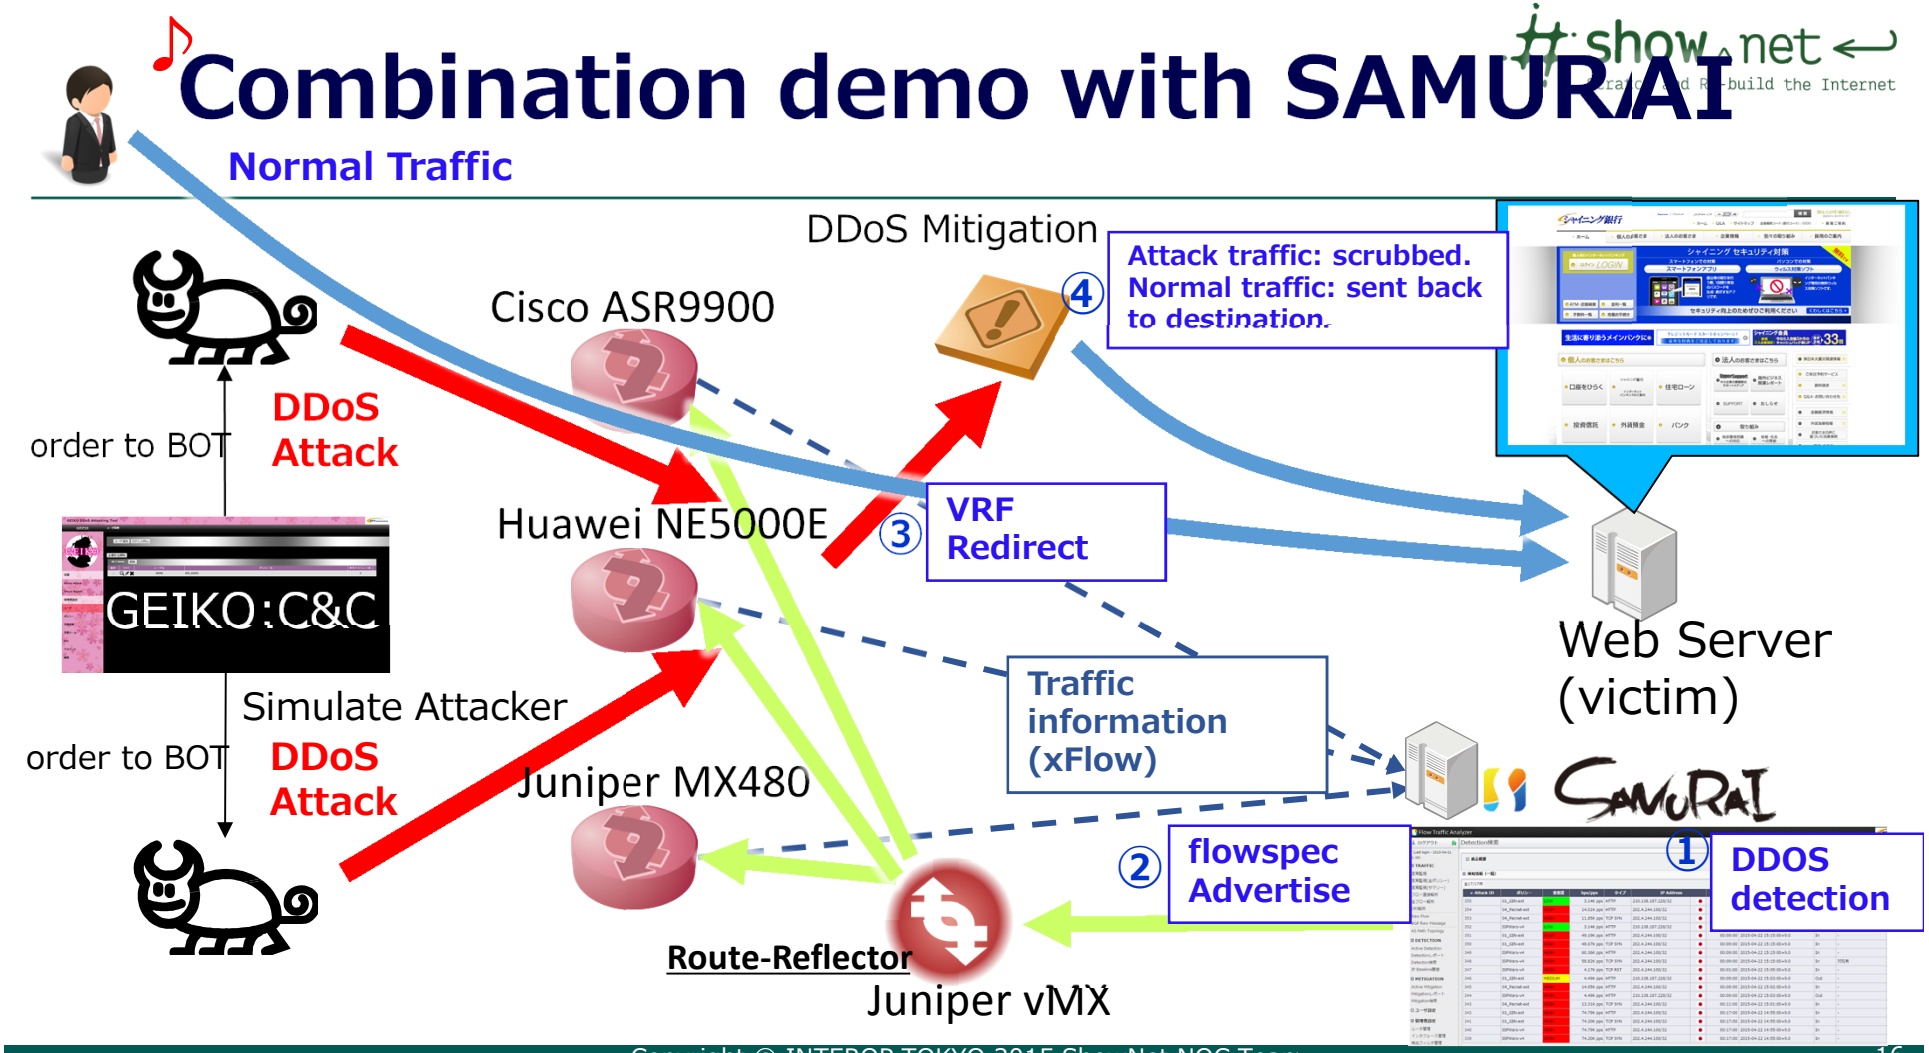

Summary

- Operator very interested in BGP flowspec
- Need more multi vendor interop report
- We confirmed 4 vendors(Cisco/Juniper/Huawei/Samurai) interoperability
- Implementation date is quite difference , therefore detail information would be needed.
- RFC5575 description sometimes heavy to understand, sample example is helpful. (m=0 is needed??)
- IETF implementation report would be welcomed.

show net
Scratch and Re-build the Internet

Special Thanks

We appreciate a lot of support

Appendix

Software Version

- Huawei NE5000E 8.65
- Cisco ASR9900 IOS-XR 5.3.1
- Juniper MX480 Junos 15.1R1.8