IDR at IETF 94
11/2/2015 at 13:00 – 15:00
Agenda

Agenda (1)

1) Chair's slides [13:00-13:05]
presenter: Chairs
2) draft-keyupate-idr-bgp-attribute-announcement.txt
presenter: Keyur Patel
time: 13:05-13:15

3) draft-wu-idr-bgp-segment-allocation-ext-00
https://datatracker.ietf.org/doc/draft-wu-idr-bgp-segment-allocation-ext/
presenter: Shunwan Zhuang
time: 13:15 - 13:25

4)draft-sreekkantiah-idr-segement-routing-te-00
presenter: Arjun Sreekantiah
https://tools.ietf.org/html/draft-sreekantiah-idr-segment-routing-te-00
Time:13:25-13:35

Agenda (2)

5) draft-litkowski-idr-flowspec-interfaceset-01.txt
Speaker: Stephane Litowski
Time: 13:35-13:40

6) draft-liang-idr-bgp-flowspec-label-01.txt.
Speaker: Jianje You
Duration: 13:40-13:50

7) draft-liang-idr-bgp-flowspec-time
presenter: Shuwan Zhuang
time: 13:50-14:00

8) Draft Name: draft-hao-idr-flowspec-nvo3-02
Speaker: Weiguo Hao
Duration: 14:00 – 14:05

Agenda (3)

9. draft-hao-idr-flowspec-redirect-tunnel-00
   Speaker: Lucy Yong
   Time: 14:00- 14:15

10. draft-vandevelde-idr-flowspec-path-redirect
    Speaker: Gunter Van De Velde
    Time: 14:15-14:25

11. Draft Name: draft-li-idr-mpls-path-programming-02
   Speaker: Zhenbin Li
   Duration: 14:25-14:35

12 . Draft Name: draft-li-idr-flowspec-rpd-01
   Speaker: Shunwan Zhuang
   Duration: 14:35- 14:45

Agenda (4)

13. JANOG operational tests of Flow Spec
    presenter: TBD
    duration: 14:45-14:55

http://www.janog.gr.jp/en/index.php?JANOG36_Meeting%2FJANOG36_Program_Contents%2Fbgpflowspec

IDR chairs status slides

Sue Hares, John Scudder

Status of drafts

- RFCs – 3 RFCS
  - AS0 (RFC7607)
  - Error handling (RFC 7606)
  - flowspec-redirect-rt (RFC7674)
- RFC editor – 2
  - draft-ietf-as-migration,
  - Draft-ietf-ls-distribution
- At IESG:
  - draft-ietf-idr-ix-bgp-route-server
  - draft-ietf-idr-ix-bgp-sla-exchange
  - draft-ietf-rtc-no-rtc
- Early allocation
  - draft-ietf-idr-bgp-extended-messages-09
  - draft-ietf-idr-bgp-prefix-sid

Status of drafts

- new drafts
  - draft-ietf-idr-route-leak-detection-mitigation-00
  - draft-ietf-idr-tunnel-encaps-00
  - Draft-ietf-idr-bgp-prefix-sid
  - Draft-ietf-idr-ls-trill-00
- Work focus
  - Yang models for BGP,
  - Flow Specification+, ls-distribution+, RTC+,
  - NextHop+/tunnels+, segment routing+, prefix-si,
  - Route leak-mitigation
  - Bgp-extended-message, bgp-gr-notification
- Charter open issues – Push toward IEG
  - Add path (in progress), custom decision
  - ASPATH ORF, Multi-session BGP, Dynamic capabilities,
  - BGP MIB-v2, BGP Link-Bandwdith Community
- Recharter discussion in 2016

Constrain Attribute announcement within BGP

draft-keyupate-idr-bgp-attribute-announcement-00

Keyur Patel, Jim Uttaro, Bruno Decraene, Wim Henderickx

IETF 94, November 2015, Yokohama, Japan

Motivation

- Currently there is no mechanism to scope the announcements of optional attributes

- The only possible way to filter attributes within BGP are:
  - Unrecognized Optional non-transitive attributes
  - Error handling filters malformed attributes
  - Attribute Specific rules to ensure their scope (Local Pref)

- Need for scoping attributes (atleast) at:
  - Confed boundary
  - AS boundary
  - At Multi-AS administration boundary

Use Case

- BGP Tunnel Encap attribute
  - Defined in ietf-idr-tunnel-encaps
  - Scope the Tunnel attribute announcements
- BGP Nexthop Capabilities Attribute
  - Defined in draft-decraene-idr-next-hop-capability-01
  - Optional Non Transitive Attribute defines Nexthop's capabilities
- BGP Timestamp Attribute
  - Defined in draft-litkowski-idr-bgp-timestamp-02
  - Carries Timestamps for a given NLRI for each BGP speaker the NLRI traverses
- Any new attributes defined in future…..

Solution

- No use of Capability
  - Adds complexity to protocol
- Define 2 unused bits of Attribute flags:
  - O Optional or a Well-known as defined in [RFC4271] 1st bit
  - T Transitive or Non-Transtive as defined in [RFC4271] 2nd bit
  - P Partial as defined in [RFC4271] 3rd bit
  - E Extended Length type as defined in [RFC4271] 4th bit
  - A AS Wide Scope 5th bit
  - C Member-AS in Confederation Scope 6th bit
  - M Multi-AS Scope 5th and 6th bit
- In order to preserve the bits Multi-AS scope is enabled when 5th and 6th bits are both turned on!

Solution - Rules

- A, C OR M Bits require O bit to be set
- Filtering based on bits must be enforced when a BGP speaker receives or originates a route
- Requires implementation to enforce Enhance Error handling rules for attributes
  - Malformed attributes having impact on route selection or route installation should enforce “treat-as-withdraw” procedure
  - Other Malformed attributes should enforce “attribute-discard” procedure

Alternate Solution 1

- Reserve first 4 bytes of attribute data field for all newly allocated attributes
  - Mark them as flags field
- Defined the scope bits from the reserved flag fields
- Reserve IANA space for new attributes so that implementations modify the attribute code to reserve first 4 bytes as flags field

Only makes sense if more scoping modes are needed

Alternate Solution 2

- Define new attribute for scoping attributes
- Attribute consist of one or more TLVs
  - TLV contains, Attribute type value and its scope
- Modify the code to setup the dependency for attributes

Sets up a dependency with actual attributes! Complicates the code!

Questions?

BGP Extensions for Segment Allocation

draft-wu-idr-bgp-segment-allocation-ext-00

Shunwan Zhuang(zhuangshunwan@huawei.com)
Nan Wu(eric.wu@huawei.com)

IETF94, Yokohama

Motivation (1)

☐Allocating SIDs in BGP-only Networks

- No IGP flooding mechanism to advertise SRGB.
- Be better to allocate SIDs through BGP in a centralized way.

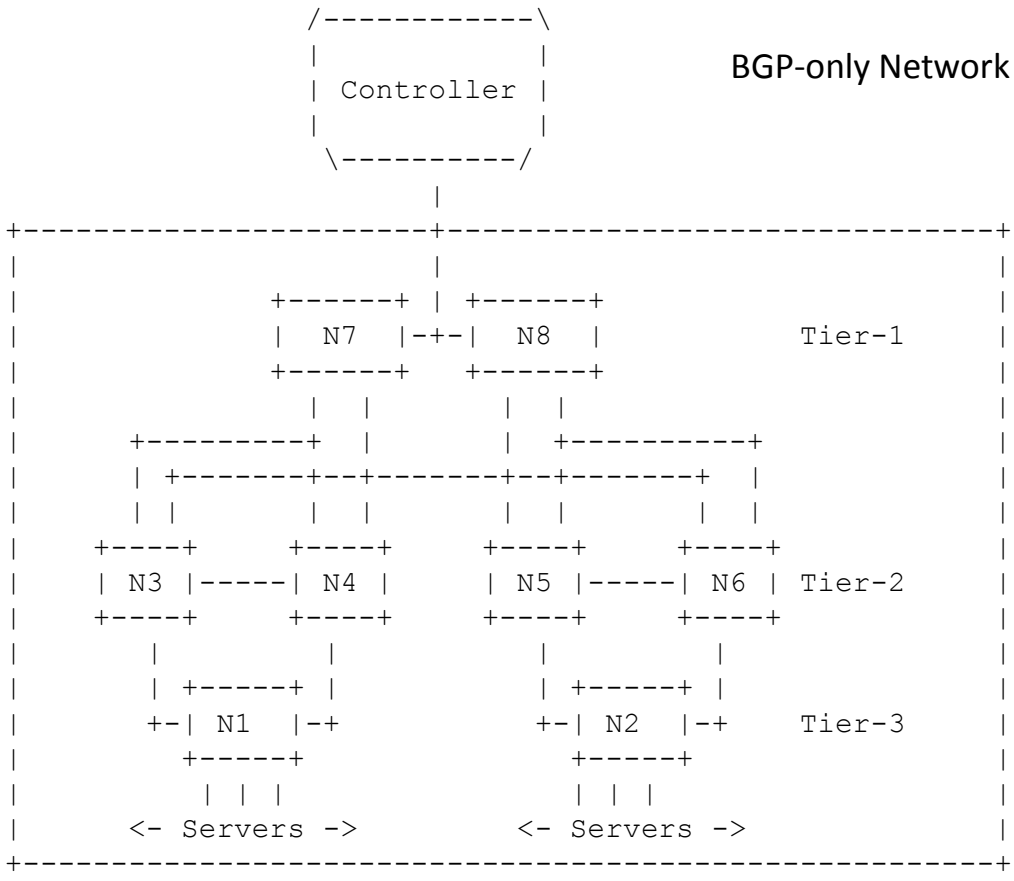

Motivation (2)

☐ Allocating SIDs to a on-site SRMS

- In SR&LDP interoperation scenario, an on-site SRMS is used.
- SIDs for mapping entries are allocated by a remote network planning tool.
- BGP may be one of candidate signaling protocols.

On-Site SRMS Case

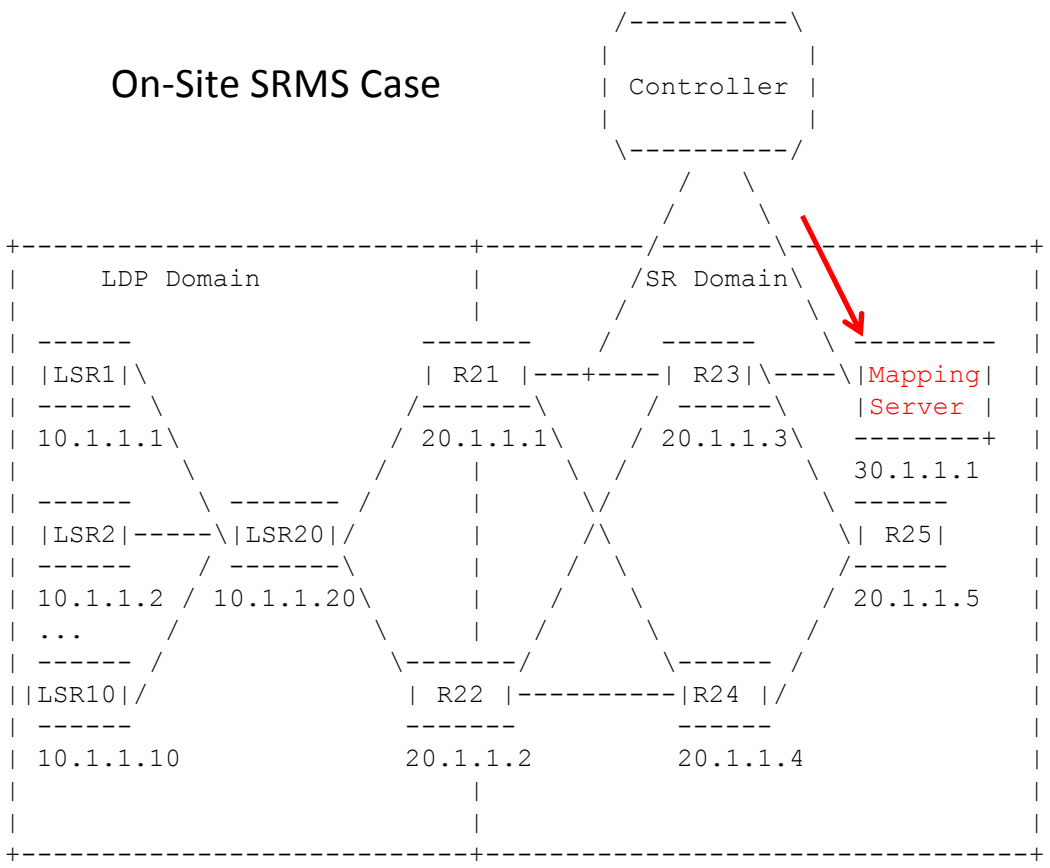

Choice of Protocol Extensions

- Option 1: Totally reuse the existing BGP-LS for Segment Allocation
  - Existing BGP-LS is always to carry IGP link-state information from IGP network to collector.
- Option 2: New BGP Extensions for Segment Allocations
  - Introduce a new protocol ID: The use of a new Protocol-ID allows separation and differentiation between the NLRIs carrying Segment Allocation information from the NLRIs carrying IGP link-state information
- The document prefers the option 2.

Protocol Extensions

- New Protocol-ID: BGP-Segment-Allocation (TBD)
- Use existing BGP-LS-TLV:

NLRI	TLV Code Point	Description
NODE	1034	SR Capabilities
NODE	1035	SR Algorithm
LINK	1099	Adj-SID
LINK	1100	LAN-Adj-SID
LINK	1036	Peer-SID
LINK	1037	Peer-Set-SID
PREFIX	1158	Prefix SID
PREFIX	1033	SID/Label Binding

Next step

- Solicit comments on the choice of protocol extensions.
- Refine this draft

A plethora of flowspec features and an afternoon chat at IETF 94

Jeffrey Haas <jhaas@juniper.net>

The current landscape

- Flowspec has become a very popular feature and is getting a significant number of additional mechanisms proposed for it. These mechanisms fall into two broad categories:
  - Additional traffic redirection behaviors.
  - Additional match criteria.
- Several, but not all, of the authors of various flowspec features met on Sunday of IETF-94 to discuss reconciling some of the potential feature interactions. This is my summary of that discussion.

Redirection (1)

- The original feature for redirection was “redirect to VRF”.
- A “redirect to IP” feature is an adopted I-D in IETF.
  - A significant portion of the complexity in the draft was reconciling behavior when both redirect to VRF is specified along with redirect to IP.
  - Additionally, the behavior when multiple redirect-to-IP extended communities was addressed.

Redirection (2)

- The new redirection functionalities fall into three broad classes with relation to redirect-to-VRF/IP:
  - Complementary: Tunnel
  - Partially complementary: MPLS Label
  - Disjoint: Path-id (?), SID

Match Criteria

- NVO3 – Applicable in the edge scenario, but not likely in transit.
- Time – Clearly specified. My specific concern is “surprises”, partially due to scenarios relating to how flowspec orders rules along with multiple distributors.
- Interface-set. Calls out potential need for add-paths support.

Informational

- The proposal adds a descriptive piece of text to the flowspec rule. Discussion has suggested if we want such a feature, it is probably a general purpose feature.

Future steps

- For features that may impact forwarding:
  - Determine if they are intended to work with each other and what the behaviors are if they do. Our challenge is where we document the full set of combinations
  - When the features are not complementary, how do we tie-break on routers that implement both features?
    - Ignore all?
    - Preference? Where is this preference? In spec, or via configuration? What about tie-breaking?

draft-litkowski-idr-flowspec-interfaceset-02

S. Litkowski, Orange
A. Simpson, ALU
K. Patel, Cisco
J. Haas, Juniper

Problem statement

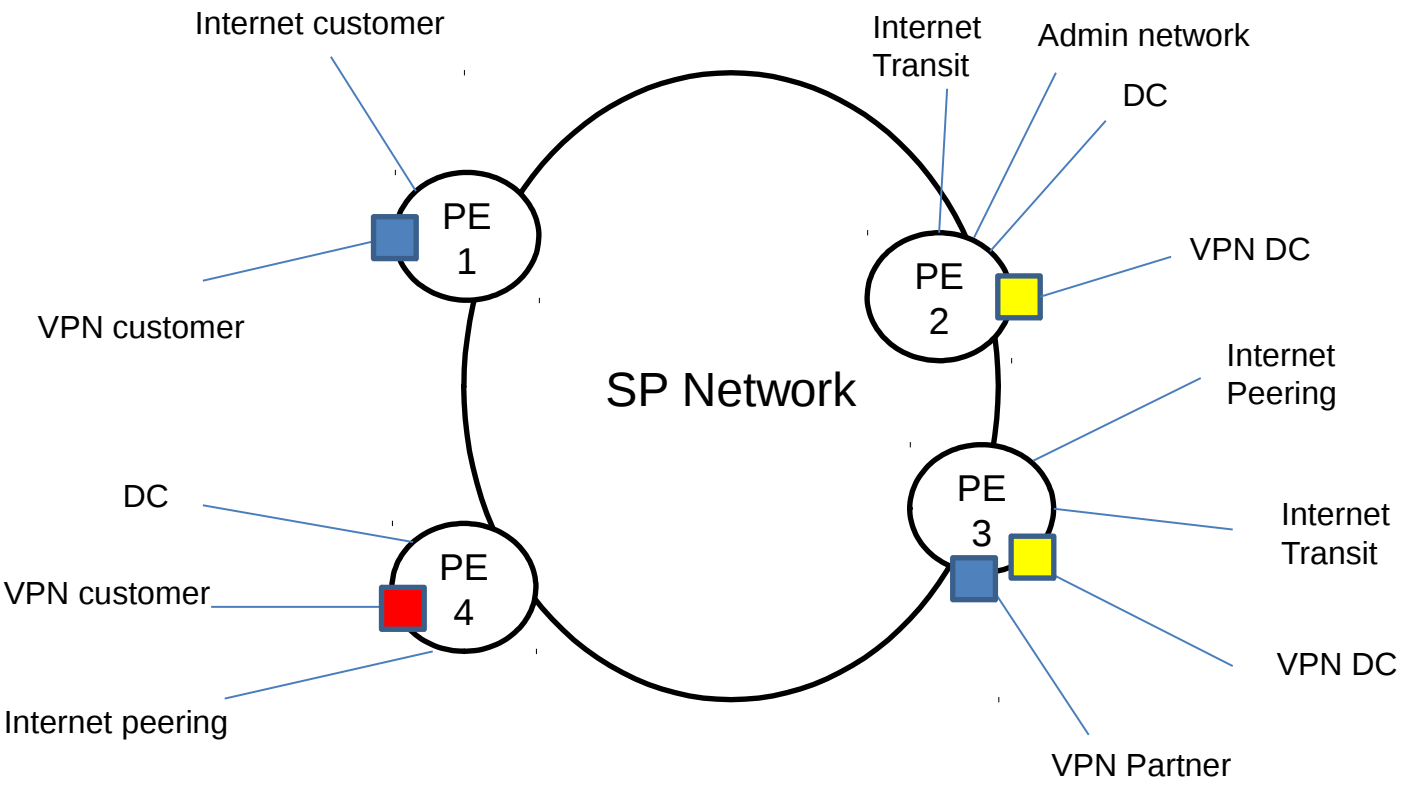

Multiple outside connections in the network

How to deploy specific Flowspec rules on a specific set of interfaces ?

interface-set extended community

- Transitive 4-B AS-specific extended community
  - Global admin : ASN of the originating router
  - Local admin :

Flags to manage direction of the Fspec rule (Input/output)

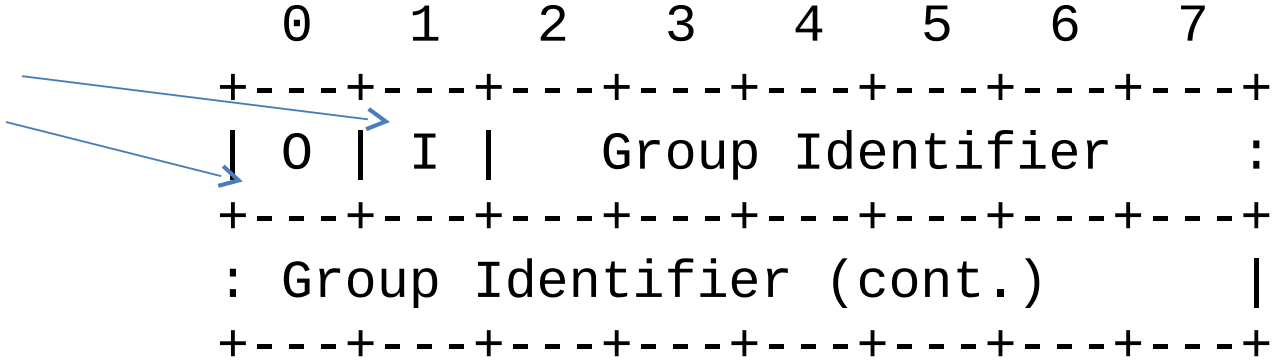

- Multiple interface-set on the same Fspec NLRI means « match-any »

WG discussions

- new NLRI instead vs EXTCT usage :
  - New NLRI may bring complexity in filtering at AS boundary (there is a high chance for intf groups to be local to an AS)
  - Use of EXTCT requires ADDPATH in some cases which is acceptable

- AS4Bytes extct vs RT :
  - We may need to define a transitive and non transitive version (to allow for better interAS management)
  - RT would allow for constrained route distribution
  - Reuse concepts described in MDCS draft ? (include/exclude RTs)
  - Do we define a new RT type ? (as for ES-Import) => we need to encode direction

Next steps …

- Feedbacks expressed on the list shows that this extension is useful
- Need to reach consensus on encoding
- Update the draft accordingly

Label Information for BGP FlowSpec

draft-liang-idr-bgp-flowspec-label-01

Qiandeng Liang (liangqiandeng@huawei.com)
Jianjie You (youjianjie@huawei.com)
Robert Raszuk (robert@raszuk.net)
Dan Ma (danma@cisco.com)

Status of this I-D

- First presented in IETF 93, Prague meeting
  - Would “link FlowSpec to RFC3107” satisfy the requirements?

    Though FlowSpec rule could use the label(s) bound with the best-match route to the target IP in the 'redirect to IP' action, in order to differentiate FlowSpec rules, each rule needs to be assigned a unique IP address. This would consume too much IP address resources.

- The update compared to v-00
  - Label encoded in ACTIONS section of RFC5575
  - Extend the match criteria to the label within the packet header

FlowSpec Label Action

A new label-action is defined as BGP extended community value based on Section 7 of [RFC5575].

type	extended community	encoding
TBD1	label-action	MPLS tag

Label-action is described below:

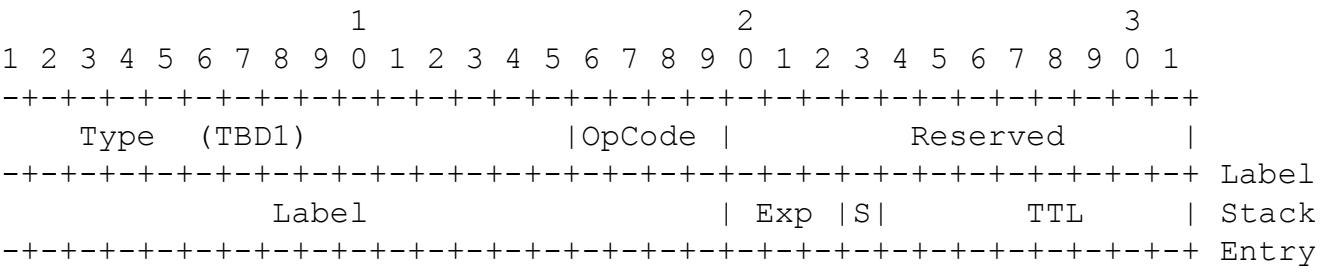

- Type: indicates the label action
- OpCode: operation code; 0: Push; 1: Pop; 2: Swap; 3-15: Reserved
- Label Stack Entry: the same as defined in RFC3032

FlowSpec Label Action

◆ If the BGP router allocates a label for a FlowSpec rule and disseminates the labeled FlowSpec rule to the upstream peers, it can use the label to match the traffic identified by the FlowSpec rule in the forwarding plane.

◆ A FlowSpec rule MAY include one or more ordering label-action(s). The arrival order of the label-actions decides the action order.

Next Step

- Accepted as WG doc?

- Solicit more comments and suggestions on the mailing list

Thank You!

Scenario

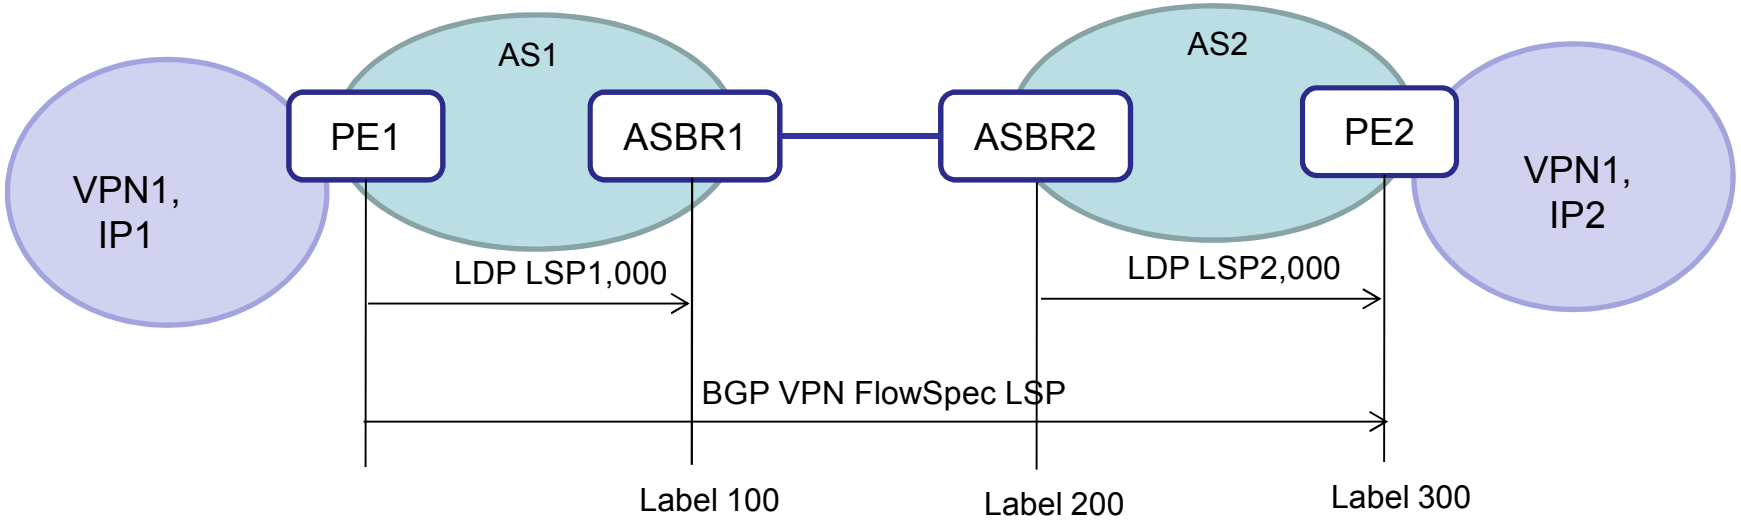

- FlowSpec Rule 1 (injected in PE2)
  - Filters: Destination IP prefix:IP2/32; Source IP prefix:IP1/32
  - Actions: traffic-marking: 1 (DSCP value)
- Forwarding Process on PE1 when receiving traffic from IP1 to IP2
  - PE1: Push 1,000 and 100
  - ASBR1: Pop 1,000, and then swap 100 to 200
  - ASBR2: swap 200 to 300, and then push 2,000
  - PE2: Pop all labels

BGP FlowSpec with Time Constraints

draft-liang-idr-bgp-flowspec-time-00

Qiandeng Liang (liangqiandeng@huawei.com)
Jianjie You (youjianjie@huawei.com)
Shunwan Zhuang (zhuangshunwan@huawei.com)

Motivation

- DDoS attacks are dynamic
  - Filtering of a flow may only be necessary for some specified time

- Deploy ACLs dynamically and flexibly
  - Steering traffic may only be necessary for some specified time.
  - Currently ACL with Time-Range is a popular feature. If we deploy ACL with Flowspec, a Time-Range descriptor needs to be considered.

FlowSpec Extended Attribute

A new BGP path attribute called “Flow Extended Attribute” is defined; it carries expected valid period information for a FlowSpec rule.

This “Flow Extended Attribute” is an optional transitive attribute that is composed of a set of Type-Length-Value (TLV) encodings, including Flow Description sub-TLV and Flow Validation Period sub-TLV.

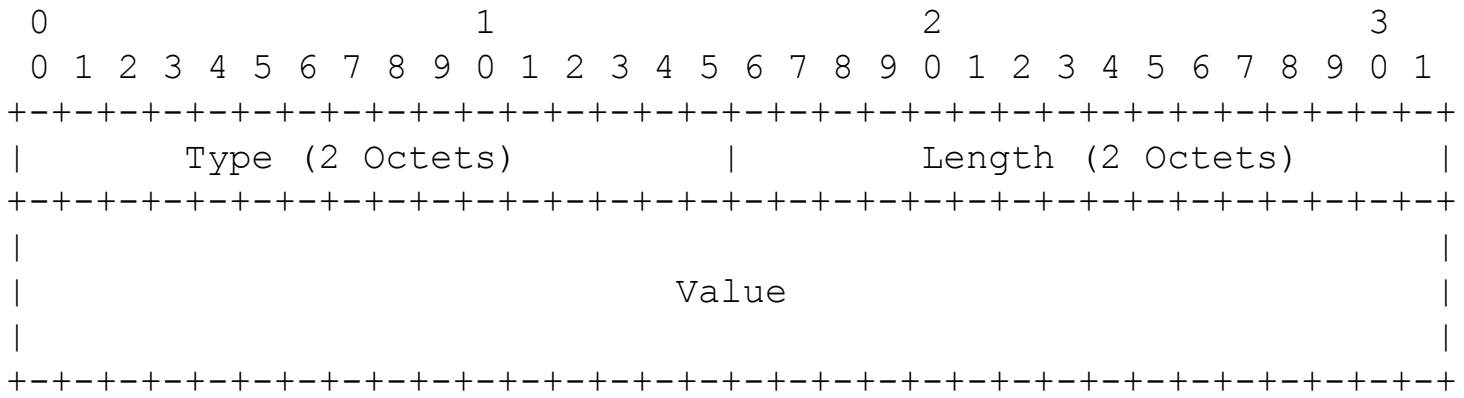

Flow Validity Period sub-TLV

The Validity Period sub-TLV is the expected valid period information for a FlowSpec rule.

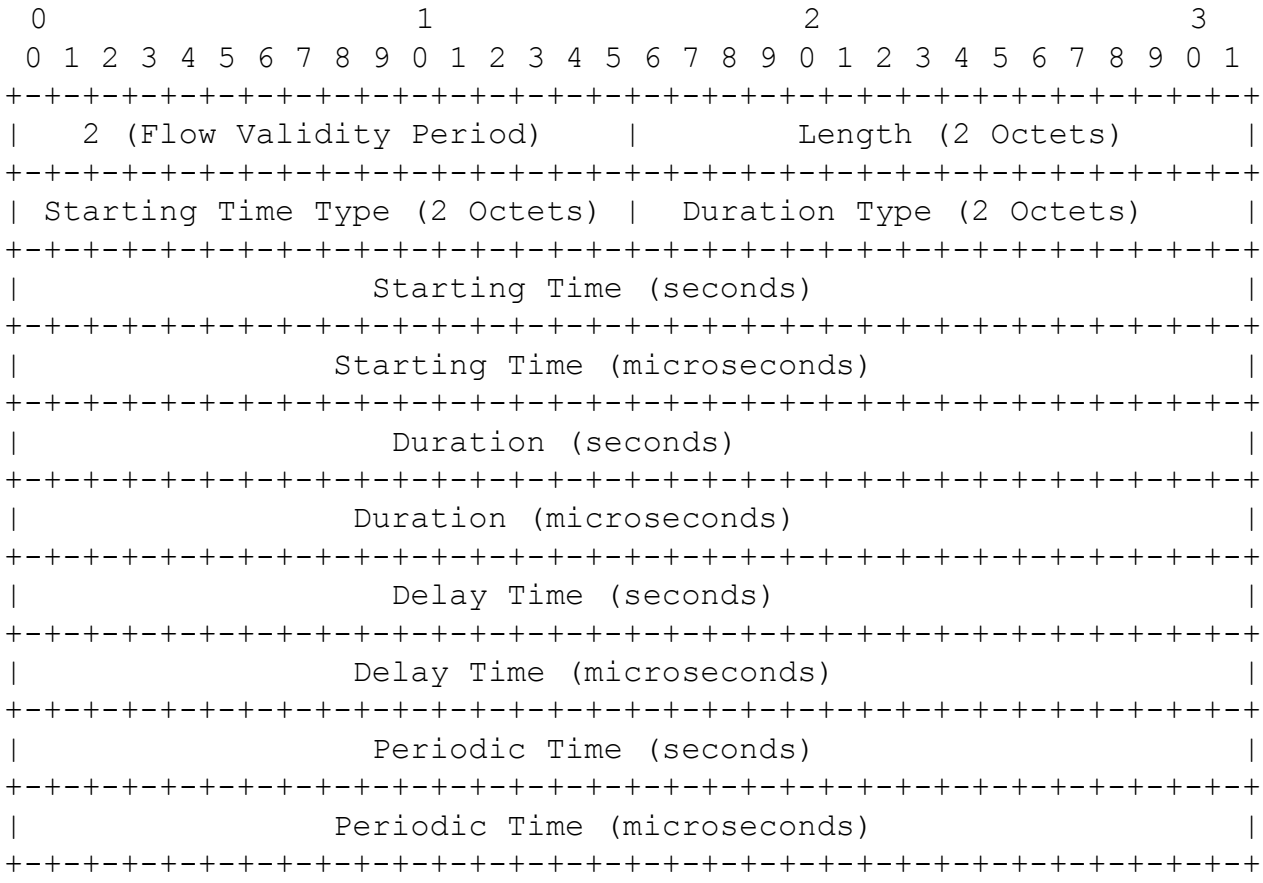

Flow Description sub-TLV

The Flow Description sub-TLV is usually used as a flow name or flow function description, in order to make the FlowSpec rule more readable in diagnosing and logging.

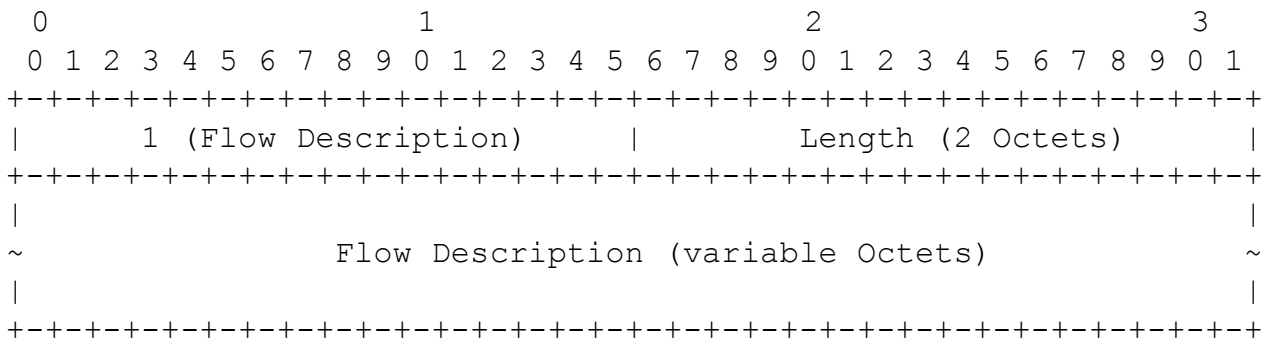

Questions from Mailing List

- Issue of usecases
  - Existing example such as ACL with Time Range.
- Should we limit it to FlowSpec AF?
  - Maybe applied to other AFI/SAFI in the future. Until now , the possible usecases confine to FlowSpec AF.
- Why choose a new path attribute rather than individual new flow components?
  - Filtering a flow through Time-Range maybe a new usecase out of the scope of this draft.

Next Step

- Sync with drafts in other WGs such as draft-zhuang-teas-scheduled-resources-00, etc.
- Solicit more comments
- Refine the draft

Thank You!

BGP Flowspec Redirect-to-tunnel

draft-hao-idr-flowspec-redirect-tunnel

Presenter: Lucy Yong lucy.yong@huawei.com

November 2015 , Yokohama Japan

IP Flow Redirect-to-tunnel

A. IP flow packets are redirected/copied to the third place (i.e. not BGP next hop) via a tunnel
   - E.g. flow packet inspection, SFC, traffic-engineering, etc

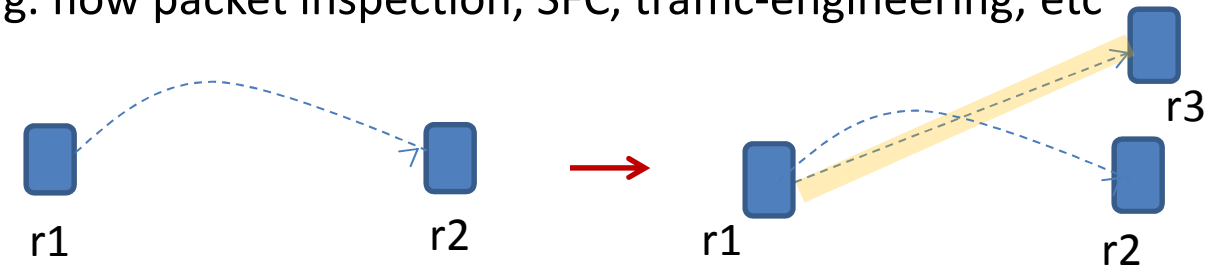

B. IP flow packets are redirected to the BGP next hop via a tunnel
   - E.g. configuration automation, differentiated services, traffic-engineering

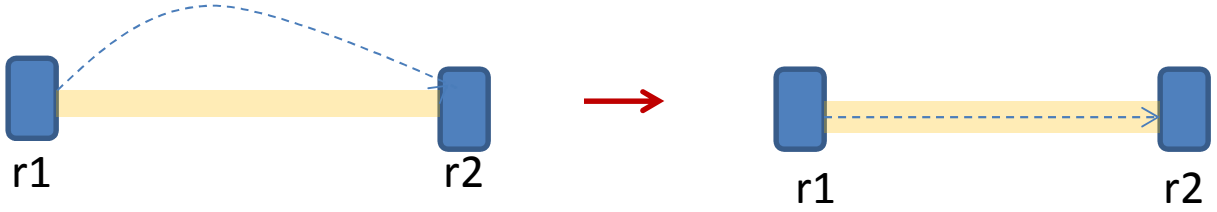

Question: What kind of tunnels make a sense to be used in case A and case B?

Two Types of Tunnel

- IP Tunnel: a tunnel whose endpoints are identified by IP address. Tunneled flow packets are carried as a regular IP packets over IP network

- Non-IP Tunnel: a tunnel that is implemented by non-IP data plane, tunneled packets are carried by non-IP data plane, e.g. MPLS data plane.
  - The tunnel is constructed by other control protocol such as RSVP-TE, LDP, etc

- Question again: will both types of tunnel apply to case A and case B?

Tunnel Policy

A. A tunnel can be identified by the tunnel type and criteria; e.g. IP tunnel, MPLS LSP (e.g. LDP RSVP-TE, SR-BE, SR-TE)

B. A tunnel is identified by a tunnel identifier such as [RFC3209]

These policies are specified in BGP attributes (draft-ietf-idr-tunnel-encap; draft-li-idr-mpls-path-programming)

BGP Extension - Flow Redirect-to-tunnel

- For IP Tunnel: tunnel encapsulation attribute (draft-ietf-idr-tunnel-encaps) is used to convey IP tunnel info. for a flow-spec,
  - Extend the attribute to SAFI 133,134
  - Specify the semantics for the SAFIs and usage constraints

- For MPLS Tunnel with tunnel type: mpls tunnel type and criteria carried by tunnel encapsulation attribute (draft-li-idr-mpls-path-programming)
  - Extend MPLS tunnel types and sub-TLV

- For MPLS tunnel with tunnel identifier: tunnel identifier carried in BGP extended unicast tunnel attribute (draft-li-idr-mpls-path-programming)

The BGP extension is used for tunnel selection for the flow spec, not for tunnel construction.

Next Steps

- Update the draft to clarify the semantics and usage constraint for a flow carried via an IP tunnel
- Clarify the semantics and usage for a flow carried via MPLS tunnel
- Solicit comments
- Work with the community for flowspec redirect integrity (redirect vrf, redirect ip, redirect tunnel)

Dissemination of Flow Specification Rules for NVO3

draft-hao-idr-flowspec-nvo3-02

Weiguo Hao Huawei

October, 2015 Yokohama

BGP Flow-spec for NVO3 Requirements Summary

① The match part should include inner L2/L3 header information and NVO3 header.

➢Currently the Flow specification rule [RFC5575] only includes single layer IP information, the match part lacks layer indicator and NVO3 header information, so it can't be used for the traffic filtering based on NVO3 header or a specified layer header directly.

➢This draft proposes a new subset of component types to support the NVO3 flow-spec application.

Flow-spec extension(1)

Add new Component Types for NVO3 header:

`Type TBD1 –` Delimiter type

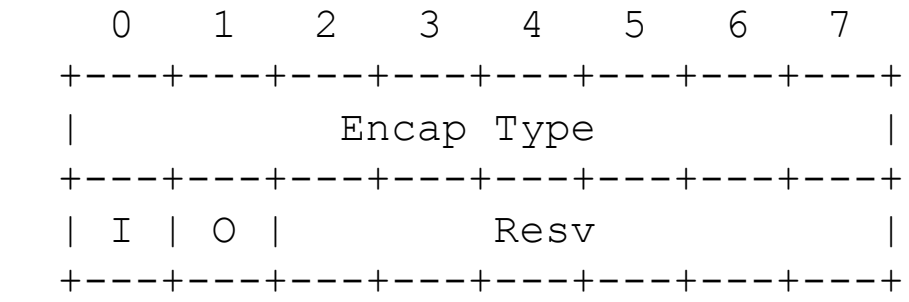

- VXLAN: Tunnel Type = 0
- NVGRE: Tunnel Type = 1

I: Inner layer indicator

O: Outer layer indicator

Flow-spec extension(2)

Add new Component Types for NVO3 header:

Type TBD2 – VNID
Type TBD3 – Flow ID, only for NVGRE
Other types:
The additional types for GENEVE [GENEVE],GUE [GUE] and GPE [GPE] header specific part will be added later.

Next Step

- Seek some comments and feedbacks
- WG adoption?

Flowspec Path-id Redirect

(draft-vandevelde-idr-flowspec-path-redirect)

Gunter Van de Velde
Wim Henderickx
Keyur Patel

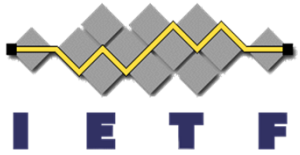

IDR Interim, 26 October 2015
Webex Virtual Meeting

Flowspec Path-id Redirect

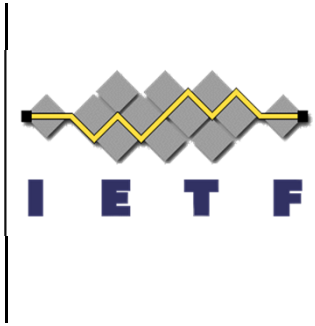

- Use-case: Traffic Steering
  - Provide a scalable apparatus to selective steer traffic onto an Tunnel (or Interface)
  - Routing system to propagate Redirect Traffic policies
- Non Use-case: Tunnel Setup signaling
  - No signaling of encapsulations
  - No signaling to setup a tunnel
  - No signaling for tunnel TE operational purpose

Anatomy of PBR

- Policy Routing has two key components
  - Identify “interesting” traffic
  - instruct what action to do with the “interesting” traffic
    - Actions
      - Traffic Conditioning
        - Policing
        - Shaping
        - DSCP/Precedence rewrite
      - Traffic Steering
        - Redirect to VRF, Interface, VLAN, next-hop, etc

Path_ID Redirect Focus

  - Note that PBR is NOT used to initiate/create/setup Tunnels

Anatomy of PBR (Cont.)

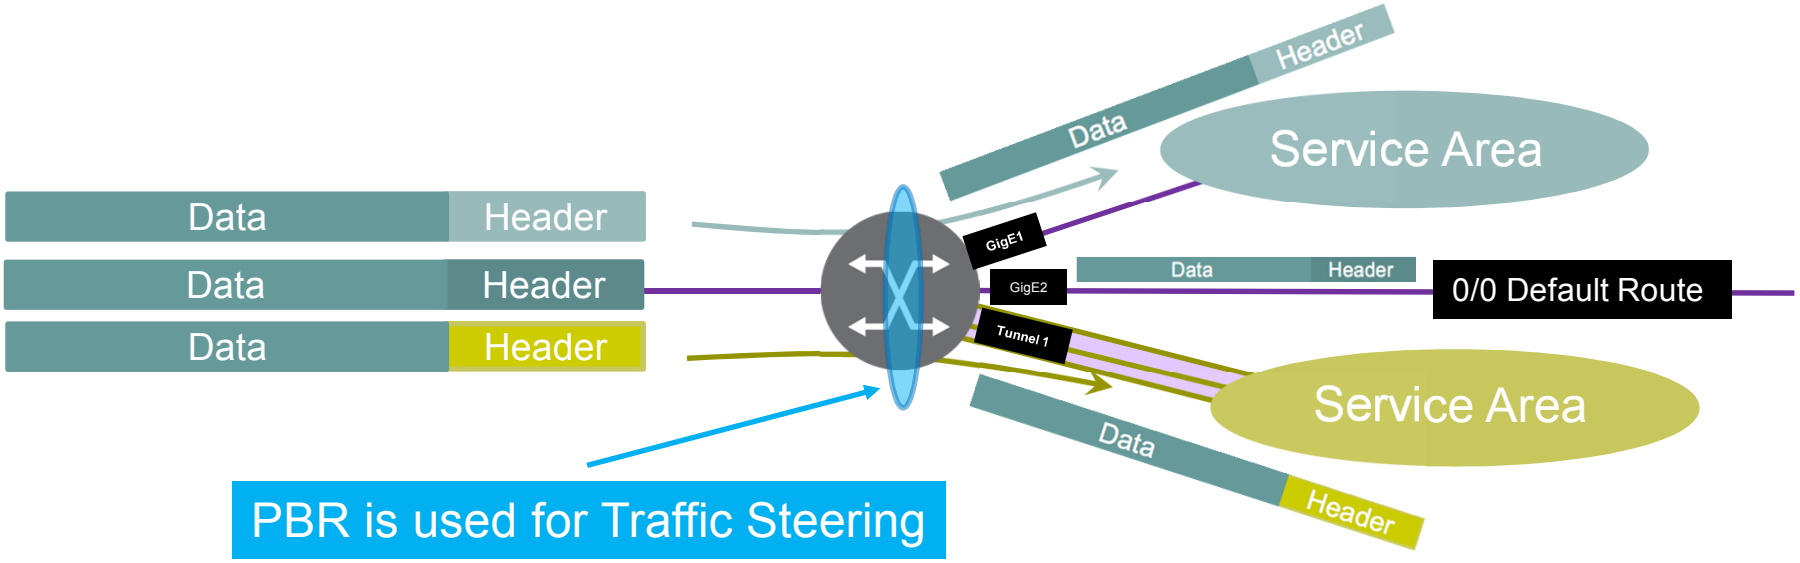

PBR rule contains Traffic Steering information

Interface, VLAN, Tunnel (RSVP-TE, SR-TE, LDP,)

Introduction to PATH_ID

- In the example traffic is redirected to different types of Interface (GigE and Tunnel)

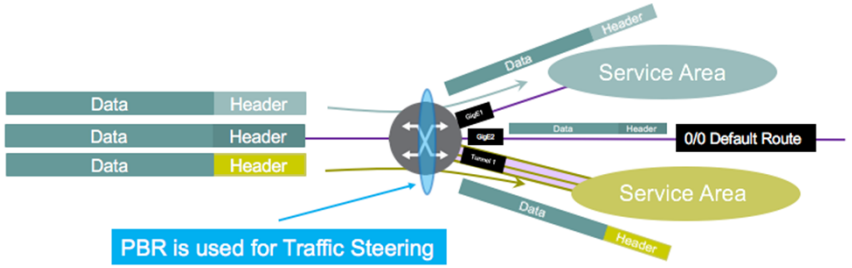

- There be dragons – there is complexity
  - Redirection interface tends to be a router local decision
  - PBR Redirection Interface to facilitate a network service most likely different per router (driven by router diversity, card-type, chassis, etc)
  - A tunnel is setup between head- and tail-end and hence uniquely signaled per router
  - When using traditional PBR, it is assumed that Tunnels are up and running before the PBR rule becomes active (This behavior should be replicated)
  - Misconfiguration of PBR could result in dramatic traffic forwarding issues

Introduction to PATH_ID (Cont.)

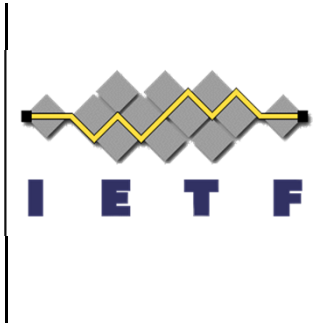

- The Noble PATH_ID Goal
  - Have a central controller send out a single unique network wide redirection policy
- Question: Can it be done by a central controller?

  (i.e. Controller (i.e. RR) sends network wide a single blur of Redirect information using BGP)
  - Type-1: Router localized recursion is possible
    - If redirection is an IP Next Hop or a redirect VPN then router can use localized recursion to discover the localized egress interface/encapsulation
  - Type-2: Router localized recursion is NOT possible
    - ~~Non-solution: configure all routers with same interfaces and tunnels~~😊
    - Solution: Create abstraction "PATH_ID" to have router localized recursion from a single unique network wide identifier to localized ingress interface/encapsulation
    - i.e.
      - Router receives redirection to PATH_ID#1 then traffic is redirected to GigE1
      - Router receives redirection to PATH_ID#2 then traffic is redirected to GigE2
      - Router receives redirection to PATH_ID#3 then traffic is redirected to Tunnel1
- Creation of Local PATH_ID to Interface/encap recursion Table?
  - Manual configuration
  - Use identifiers which already exist (PCE PLSP-ID, etc..)
  - Orchestration
  - Extensions to existing protocols

Introduction to PATH_ID

1. Controler send redirect policy into network
2. Recursive lookup for localized forwarding on ingress Packet1 (Header1->PATH_ID#1->GigE1)
3. Packet1 is forwarded onwards to GigE1
4. New ingress packet received and new recursive lookup done again. This time Tunnel2 is the Redirect result
5. New packet forwarded onto Tunnel2

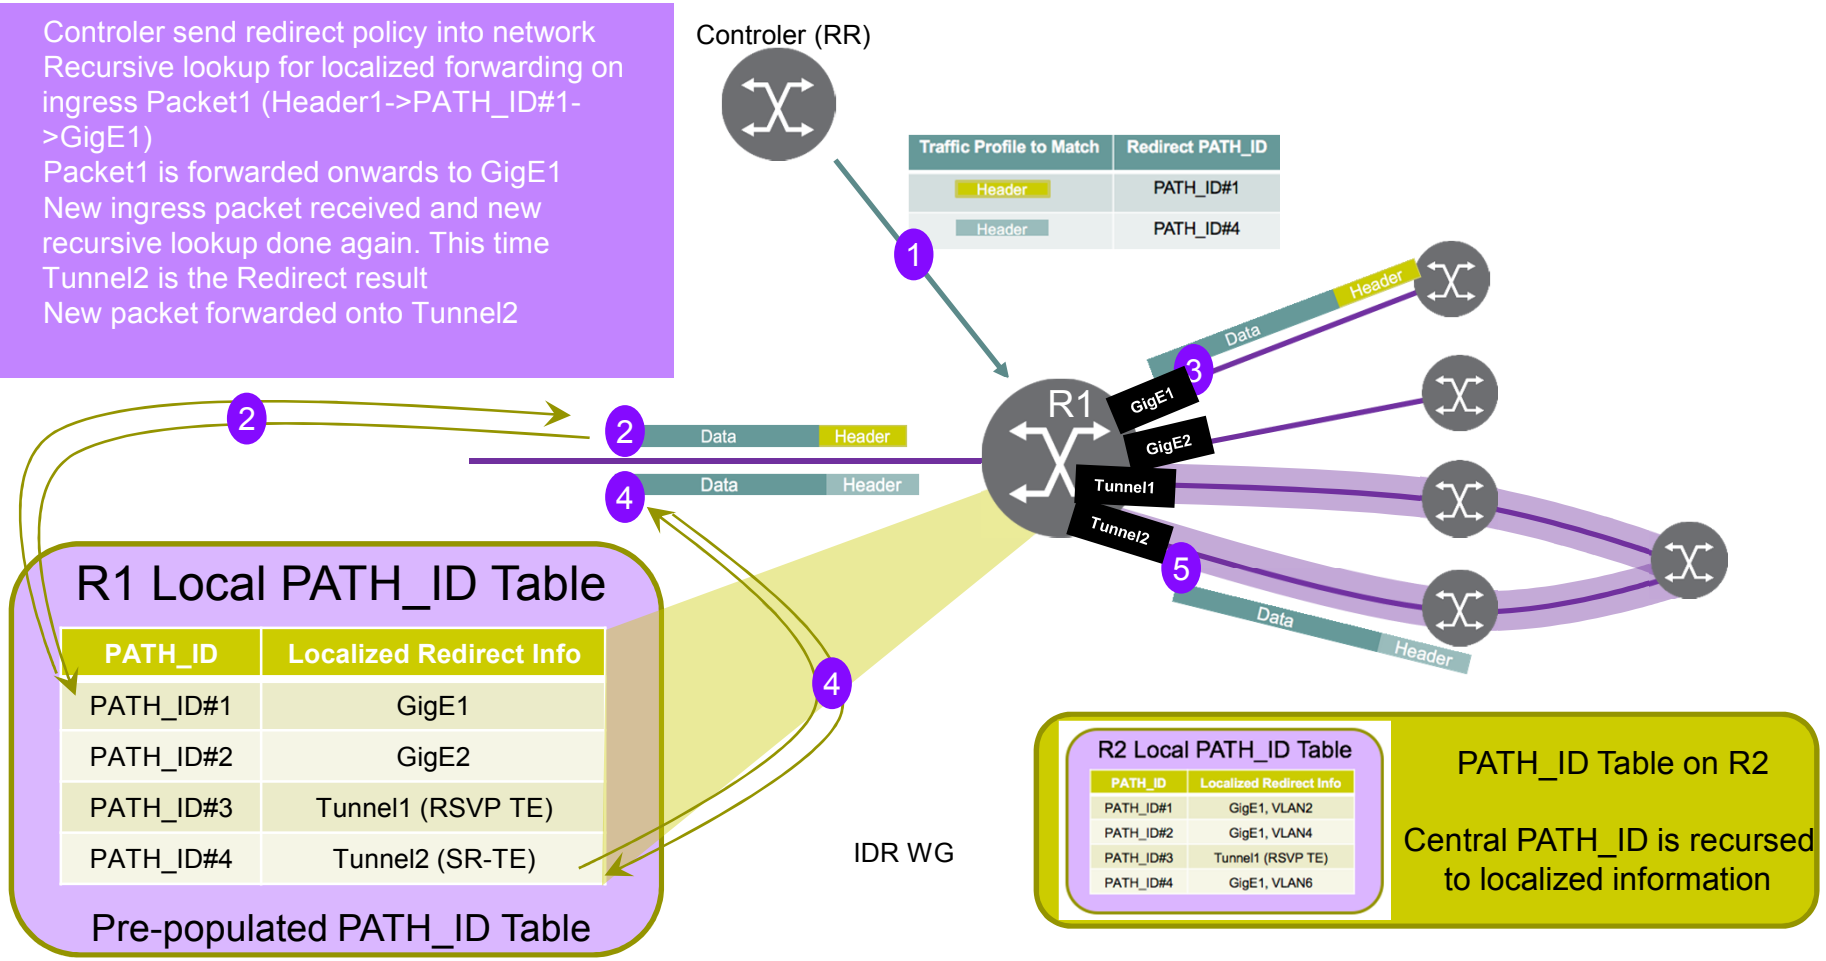

Details: Flowspec Redirect-to-PATH_ID

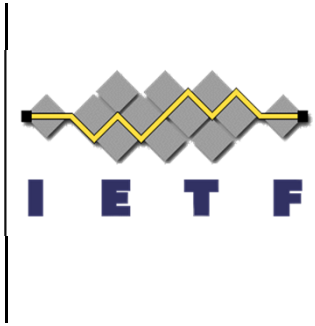

- New Flowspec Traffic Action Community
- PATH_ID is either 32 or 128 bit identifier
- Assumption
  - Router has PATH_ID table pre-populated
  - Population of this table is outside the scope of Flowspec Redirect-to-PATH_ID (work for RTGWG?)
  - Each PATH_ID is network wide unique and represents a Redirect Service identifier
- PATH_ID structure

```
0                   1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Reserved          |TID|C|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

  - PATH_ID is 32 or 128 bit value
  - C-bit (1 bit): copy original packet onto the Re-direct
  - TID (2 bit): support for nested redirects (i.e. SR Segments) or Multi-path functions
- PATH_ID decouples the Redirection Service from Redirection Interface/Encapsulation
- Note: PATH_ID could also be seen as Superset of Redirect-to-IP where Path_ID has additional context as IP address (in this case PATH_ID table and Redirect IP address are the same)

Looking at WG questions

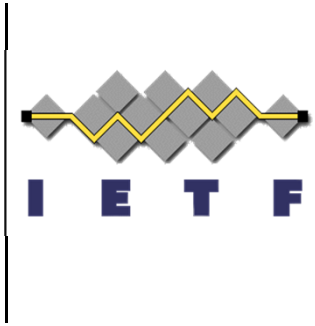

- Difference with other flowspec redirect-to-tunnel drafts?
  - This draft does not signal tunnel setup information unlike other proposals
- Purpose of TID: nested tunnels, Multi-Path, push SR segments
- If the PATH_ID is down/non-exist in the PATH_ID Table
  - If the next-hop or interface is down, then just like PBR behaviour the rule is not applied on the router. No difference with PBR behaviour from this perspective
- PATH_ID Table questions
  - Construction is outside scope
  - It could be populated CLI, Netconf/Yang, protocol extensions, etc.. (see before)
- Difference between Redirect-to-IP and Redirect-to-PATH_ID is small
  - Redirect-to-PATH_ID is indeed superset of Redirect-to-IP (in Redirect-to-ip the 32/128 bit number has IP address context corelated)
- Tunnel Setup Questions
  - Tunnel Setup is outside scope of this draft

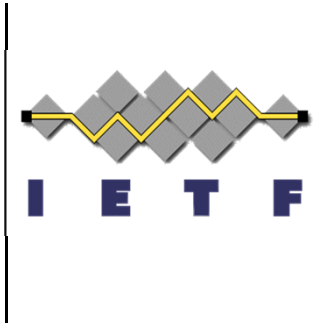

THANK YOU!

BGP Extensions for Service-Oriented MPLS Path Programming (MPP)

draft-li-idr-mpls-path-programming-02

Zhenbin Li (Presenter), Shunwan Zhuang
Huawei Technologies
Sujian Lu
Tencent

IETF 94, Yokohama, Japan

Service Oriented Segment Path Programing (SO-MPP)

- **SO-MPP - MPLS path for specific service flows**
  - Flexible label mapping to indicate service flows
  - Flow identified by 5 tuple of IP header
  - Flow mapped to MPLS Tunnel
- **BGP extensions**
  - Extended Label Attributes: Label Stack for NLRIs
  - 4 tunnel types to draft-ietf-idr-tunnel-encaps
  - Route Flag Extended Community or Cost Community
  - Extended Unicast Tunnel attribute
  - Destination Node attribute
  - Capability to signal So-MPP
- **Use Case:**
  - Deterministic ECMP on entropy label
  - Centralized controller

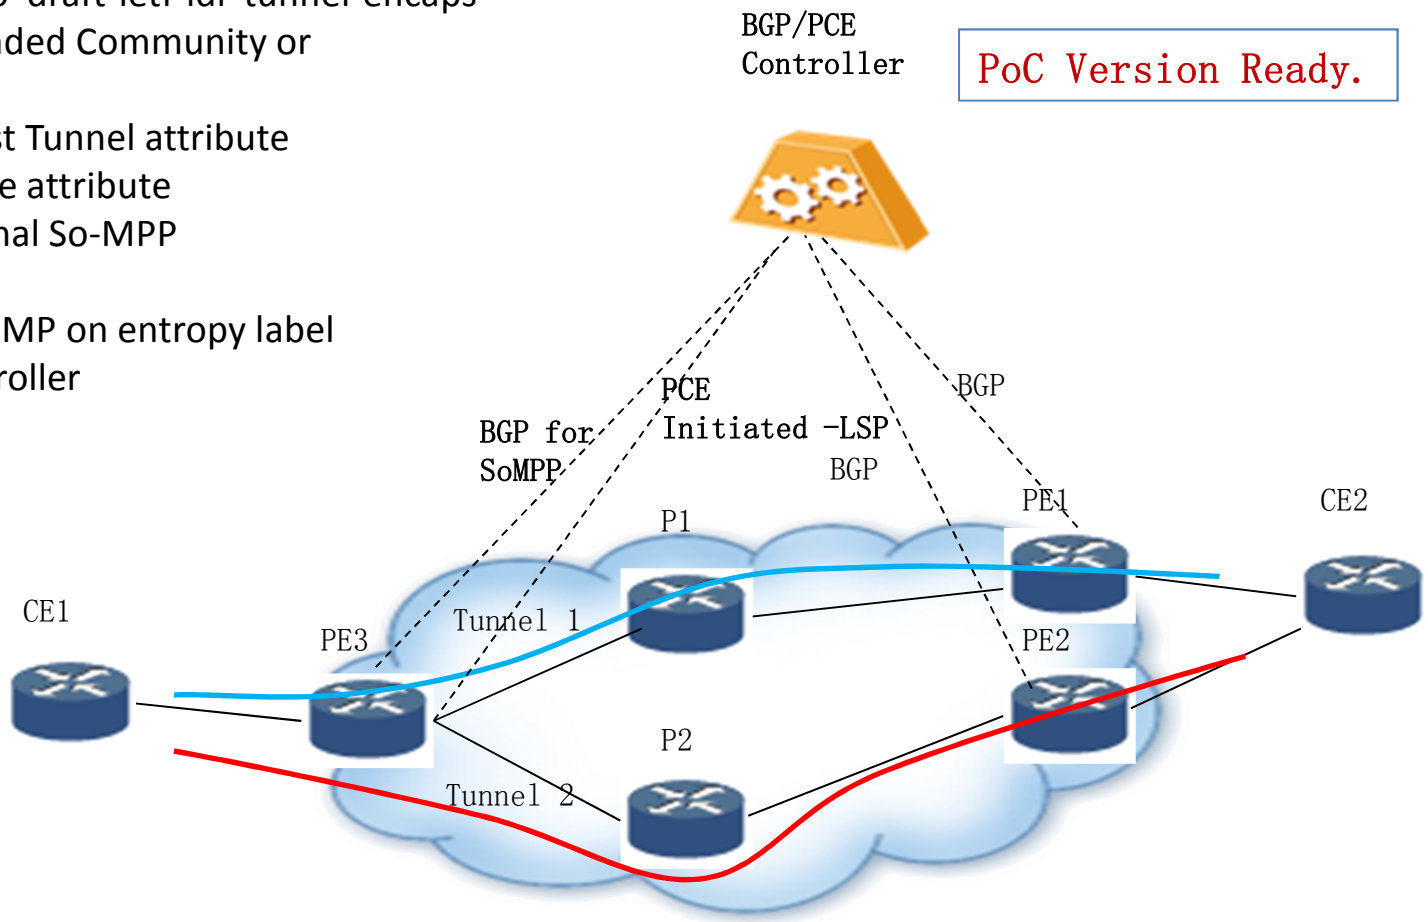

BGP Extensions (1)

- 4 tunnel types
  - LDP LSP,
  - RSVP-TE LSP
  - MPLS segment routing Best Effort,
  - MPLS-based Segment routing Traffic Engineering

- Extended Unicast Tunnel Attribute – optional transitive

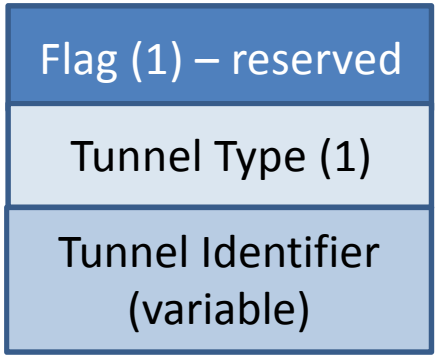

Tunnel type:
0: none (*0)
ingress
1: RSVP-TE LSP (*1)
2: MPLS segment routing TE (*1)

- Extended Label Attributes

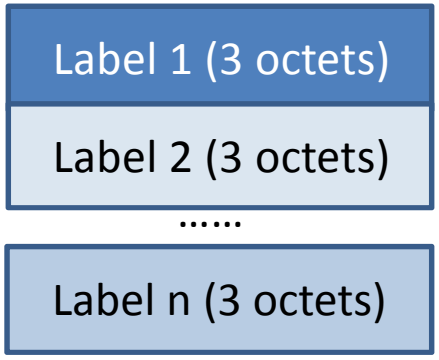

*0 – ingress router set path

*1 = <C-Type, Tunnel Sender Address, Tunnel ID, Tunnel End-point Address>

BGP Extensions (2)

- Route Flags as Best Route
  - Cost community Extended to have point insertion = 128
  - Route flag community or extended community
- Destination Node attribute – optional non-transitive

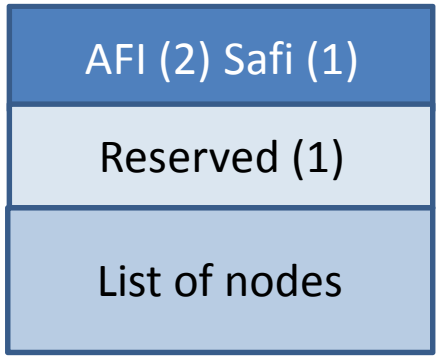

Prototypes and Open Source

- Prototypes:
  - Flexibly map services to tunnels: It has been implemented.
  - Flexible label combinations: It will be delivered soon
- Tencent has deployed the PCE/BGP controller in the network and helps refine the draft.

  Co-author (Sujian Lu (Tencent)
- Delivery of protocol extensions based on Open Source Controller is in plan.

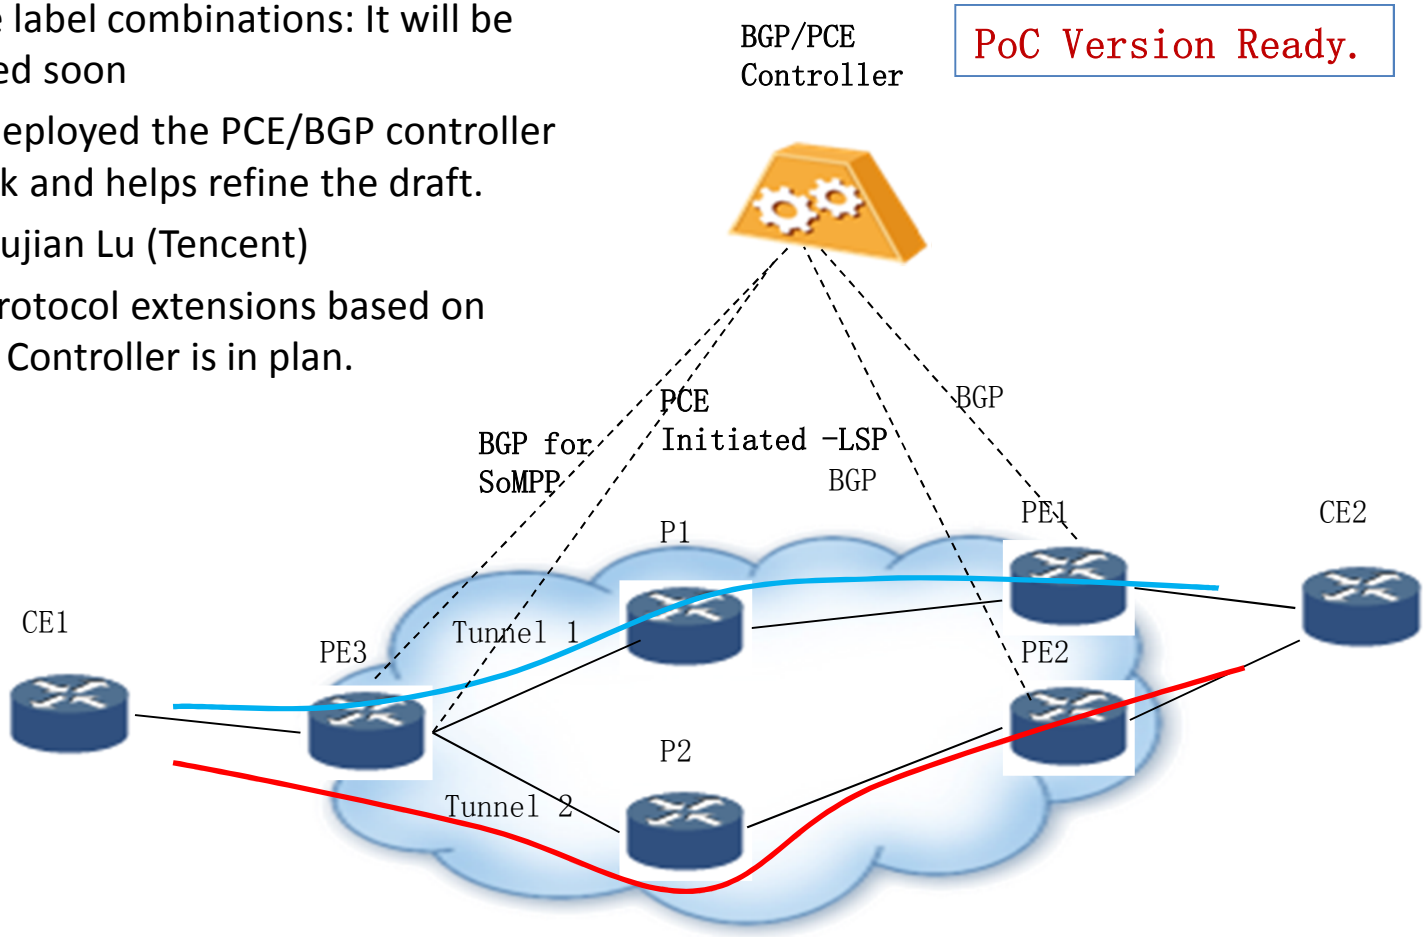

Link to Flowspec

- BGP Flowspec “Redirect to tunnel”
  - Adds tunnel types for flow
  - Utilizes draft-hao-idr-flowspec-redirect-tunnel
- BGP Path
  - draft-vandevelde-idr-flowspec-path-redirect-00
  - Path ID may specify specific tunnel.
  - Extended Unicast Tunnel may do the same thing

Next Step

- The work is practical and becoming mature now. After the refinement, the draft is also consolidated.
- Call for WG adoption.

DETAILS ON DRAFT

SPP and SoMPP

- Segment Path Programming (SPP)
  - Concept:
    - Flexible Segment Combination
    - Flexible Mapping of Service to Segment Path
  - draft-li-spring-mpls-segment-path-programming
  - Segment and Segment Path are generalized more than Segment Routing.
- Service-oriented MPLS Path Programming (SoMPP)
  - Concept: Programming MPLS path for specific service flow
    - Flexible label combination should be applied to the flow to indicate a series of service process.
    - Flexibly map the flow to tunnels with different constraints/attributes.
  - The flow can be specified by the prefix or “5-tuple” of IP header. BGP is the appropriate protocol.

Architecture and Usecases for SoMPP

- Architecture
  - Central Control for whole network & portions of network
- Usecases
  - Deterministic ECMP based on Entropy label/Flow label:
    - calculated centrally against the global view of traffic pattern
  - Centralized Mapping of Service to Tunnels:
    - PCE-initiated LSP are adopted to set up tunnels centrally with different constraints.
    - BGP extensions map the flow to the tunnel based on service requirements.

BGP Extensions for SoMPP

1. Capability Negotiation
2. Download of MPLS Path
   - Extended Label Attribute: Apply multiple labels to specific BGP prefix of multiple BGP AFI/SAFI.
3. Download of Mapping of Service Path to Transport Path
   - Specify Tunnel Type for the flow specified by BGP prefix:
   - Specify Specific Tunnel for the flow specified by BGP prefix
   - Extended Unicast Tunnel attribute is introduced to specify the tunnel identifier of RSVP-TE LSP and Segment Routing-TE path.
4. Best Route Selection:
   - Option 1: One new Extended Community, Route Flag Extended Community, is introduced.
   - Option 2: Reuse the Cost Community defined by [I-D.ietf-idr-custom-decision].
5. Specify receivers of the routes advertised by controller
   - Destination Node Attribute: a list of receiving node addresses can be defined.

Updates Since IETF 93

- Add one co-author: Sujian Lu (Tencent)
  - Tencent has deployed the PCE/BGP controller in the network and helps refine the draft.
- Add the option 2 for “Best route selection” besides introducing new Extended Community, Route Flag Extended Community.
- Procedures of load balance with multiple unicast tunnels are defined explicitly.

- Remove the early ideas of MPLS service label/segment allocation based on BGP extensions.
  - draft-wu-idr-segment-allocation will explore the possible solutions.
- Remove the “Extended Multicast Tunnel Attributes”
  - Since tunnel attributes of the existing BGP-based multicast service such as MVPN/VPLS multicast/EVPN can be reused.
  - Re-defined when it is applied in the central control environment in the future.

Relationship with BGP Flowspec (1)

- Relationship with BGP Flowspec “Redirect to tunnel”
  - Specify the tunnel type: [I-D.ietf-idr-tunnel-encaps] is necessary.

  draft-li-idr-mpls-path-programming extends [I-D.ietf-idr-tunnel-encaps] to define more tunnel types.
  - Specify the tunnel identifier:
    - Until now only RSVP-TE LSP and SR-TE Path define the tunnel identifier. draft-li-idr-mpls-path-programming defines them for the “Extended Unicast Tunnel Attributes”.
    - draft-chen-pce-pce-initiated-ip-tunnel is to introduce the tunnel identifier for IP tunnels which is composed by tunnel type, source address, destination address and tunnel ID. Maybe later more Tunnel Identifiers will be defined for the “Extended Unicast Tunnel Attributes”.
  - MPLS Tunnel Types and Extended Unicast Tunnel Attributes can be applied to more BGP AFI/SAFI that BGP Flowspec.
- Remaining Issues:
  - Specify the attributes of MPLS TE tunnels:
    - [I-D.ietf-idr-tunnel-encaps] can specify the tunnel attributes. There are rich TE attributes for RSVP-TE LSP. Can it be defined in BGP extensions?
      - Same work has ever been done. But it is dropped.
      - Tunnel attributes defined by RFC 6514 does not take it into account for P2MP MPLS TE tunnel.

Relationship with BGP Flowspec (2)

- Relationship with BGP Flowspec “Redirect to Path ID”
  - Path ID may specify specific tunnel.
    - Semantics Independent: may specify the link, netxhop, tunnel.
    - The path ID should be explained locally and not defined yet.
  - Extended Unicast Tunnel Attributes can also specify specific tunnel.
    - Semantics Dependent: Only specify tunnels.
    - Tunnel identifier has global meaning which can be understood by other nodes more than the endpoints of the tunnels. It has been defined by RFC3209.
  - Different methods for BGP Flowspec to implement “redirect to tunnel”.

Additional Points from Segment Path Programming (1)

- draft-li-spring-segment-path-programming generalizes more use cases based on segment and proposes the concept of Segment Path Programming.
- Segment ID can be an indicator which are only used in the control plane other than combining with MPLS forwarding plane and IPv6 forwarding plane for segment routing.
- Segment ID can be an local/global indicator to be seen outside for the purpose of traffic steering. It can represent the link, the node, the forwarding agency, the tunnel, etc.
- Please refer to “5. Usecases of Segment Path Programming” of draft-li-spring-segment-path-programming
  - 5.3. Steering Traffic without Mapping Segment to Label
  - 5.4. Centralized Mapping Service to Tunnels

Additional Points from Segment Path Programming (2)

- The draft of “BGP Flowspec of Redirect to Segment ID” is in process to satisfy the protocol extensions requirements of draft-li-spring-segment-path-programming:
  - Segment (Stack) Identifier Attributes can be defined to indicate the entities which the traffic will be steered to. The attributes can be applied to multiple BGP AFI/SAFI.
  - One new extended community, “Redirect to Segment ID”, can be defined for BGP Flowspec.
- If the work can be done, it may propose the alternative solutions for “redirect to interface group” or “redirect to path ID”.
  - Segment ID can be the alternative way of Path ID proposed by the draft “BGP Flowspec Redirect to Path ID”. Then it may save the work to define the Path ID from the beginning. But there is some difference between Segment ID and Path ID. For Path ID, it is totally “semantics independent” while the segment type should be specified for Segment ID which will introduce something “semantics dependent”.
  - Multiple Segment IDs which are a group of indicators of link segments of a specific node can be carried with BGP Prefix which can redirect the flow to the specified interface group specified by multiple segment IDs.

Additional Points from Segment Path Programming (3)

- How to process the overlapped the solutions for the similar requirements:
  - Keep the existing method of BGP Flowspec to define the “semantics dependent” entity such as VRF, Remote IP, Tunnel, etc. ?
  - The methods to define “semantics independent” entity such as Path ID, Segment ID, Interface Group ID, etc should be consolidated. ?

BGP FlowSpec extensions for Routing Policy Distribution(RPD)

draft-li-idr-flowspec-rpd-01

Zhenbin Li(lizhenbin@huawei.com)
Liang Ou(oul@gsta.com)
Yujia Luo(luoyuj@gsta.com)
Sujian Lu(jasonlu@tencent.com)
Shunwan Zhuang(zhuangshunwan@huawei.com)
Nan Wu(eric.wu@huawei.com)

IETF94, Yokohama

Motivation

- Provider's requirements for traffic adjustment:

- Business development or network failure introduces link congestion and overload.
- Network transmission quality decreased as the result of delay, loss and need to adjust traffic to other paths.
- To control OPEX and CPEX, prefer the transit provider with lower price.

Motivation

- Drawbacks using traditional routing policy:

- Device-based manual provisioning will cause configuration burden and misconfiguration.
- Complexity keeps increased gradually and difficulty to maintain.

Automatic provisioning mechanism is needed.

Solution

☐Routing Policy Distribution(RPD)

- Taking effect on control plane
- Impact decision on remote site

☐RPD protocol: BGP Flowspec

- Filtering rule: destination for prefix1/prefix2
- Action: R-bit introduced, more info carried in new attribute

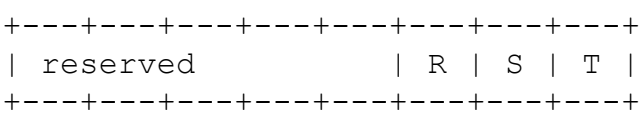

Changed from 00 version

- Alternate protocol extensions using enhanced Wide Community
- One more operator, Tencent, has similar requirements and joined in. Maybe adding new use cases in next version.

RPD Mechanism in Summary

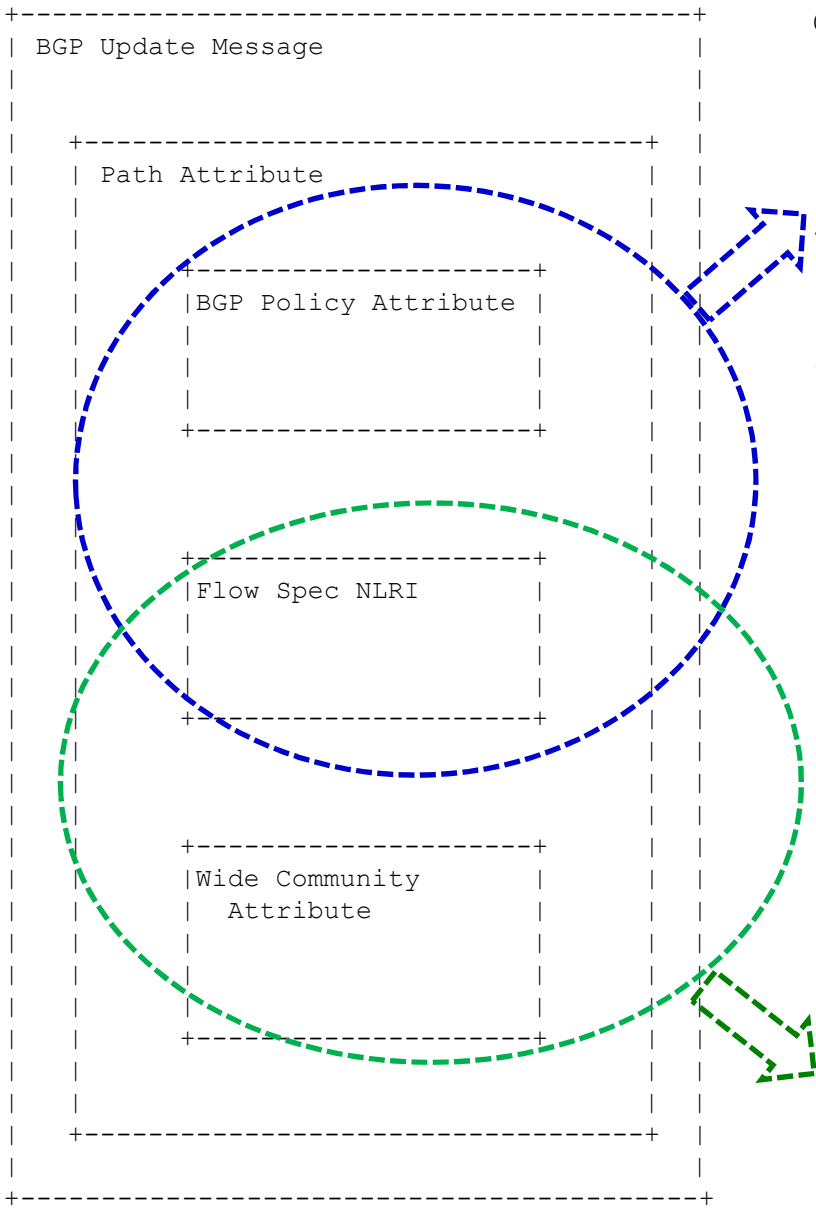

Option I:

1. Effective on which routes → Filtered by Flowspec NLRI
2. Effective on which peers → Filtered by BGP Policy Attribute
3. Take the action in BGP Policy Attribute

Option II:

1. Effective on which routes → Filtered by Flowspec NLRI
2. Effective on which peers → Filtered by Wide Community Attribute
3. Take the action in Wide Community Attribute

Protocol extensions option I(v00)

☐BGP Policy Attribute

•Attribute structure

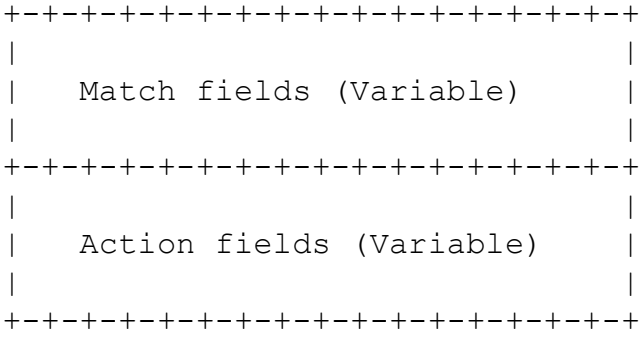

•Action field

•Action type 1: Route-Preference

•Action type 2: Route-Prepend-AS

•Match field

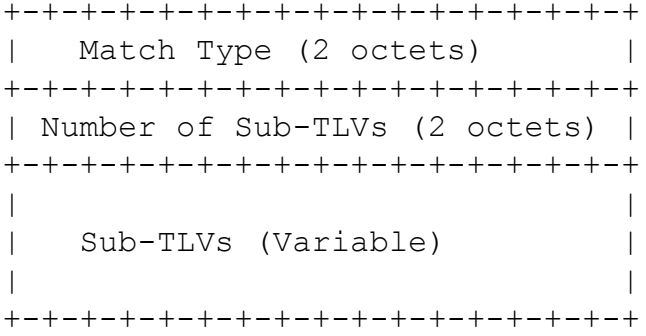

☐Match type

•Value 0: Permit, specifies the permit mode of a match rule

•Value 1: Deny, specifies the deny mode of a match rule.

☐Sub-TLVs

•Type 1: IPv4 Neighbor

•Type 2: IPv6 Neighbor

•Type 3: ASN list

Protocol extensions option II(v01)

- Wide Community is enhanced to filter a set of target routes to apply actions other than act as the attributes of advertised routes.
- New Wide Community Atoms

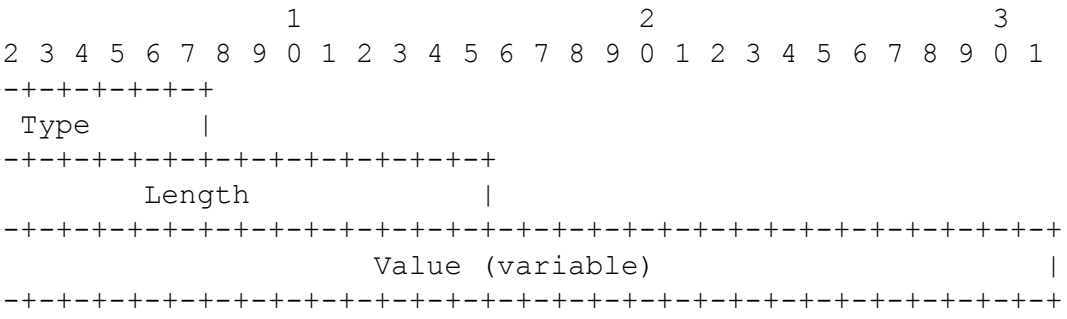

- Type 1: Autonomous System number list
- Type 2: IPv4 prefix (1 octet prefix length + prefix) list
- Type 3: IPv6 prefix (1 octet prefix length + prefix) list
- Type 4: Integer list
- Type 5: IEEE Floating Point Number list
- Type 6: Neighbor Class list
- Type 7: User-defined Class list7
- Type 8: UTF-8 String
- Type TBD: BGP IPv4 neighbor --- Newly introduced in this draft
- Type TBD: BGP IPv6 neighbor --- Newly introduced in this draft

- Actions of Wide Community can be reused and maybe enhanced in the future.

Application (1)

Inbound traffic control

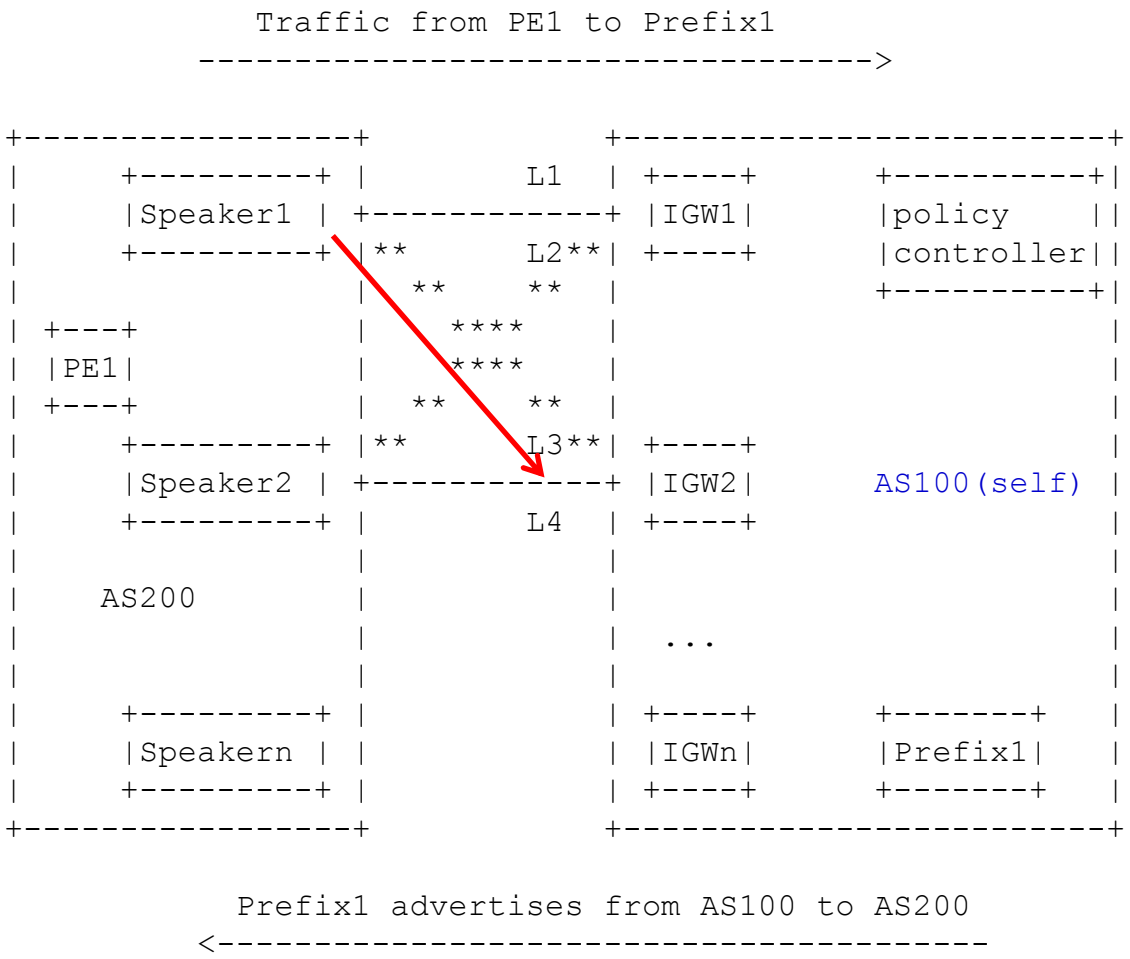

- EBGP peering:
  - Speaker1---L1---IGW1
  - Speaker2---L2---IGW1
  - Speaker1---L3---IGW2
  - Speaker2---L4---IGW2

- Requirement:
  - Administration only on AS100
  - Traffic enter AS100 through L3

→ Traffic Direction

Encoding Example (1)

☐Inbound Traffic Control encoding example

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Container Type 1 (1)          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|1 0 0 0 0 0 0 0|
+-+-+-+-+-+-+-+-+
| Hop Count: 0  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Length: 36                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Community: PREPEND N TIMES TO AS
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Own ASN                                                       100 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Context ASN#                                                  100 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|ExcTargetTLV(2)| Length:                                    11 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IPv4Neig(TBD)| Length:                                     8 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Local Speaker                                           #IGW2 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Remote Speaker                                        #Speaker1 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Param TLV (3) | Length:                                     7 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Integer (4) | Length:                                       4 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Prepend #                                                     5 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

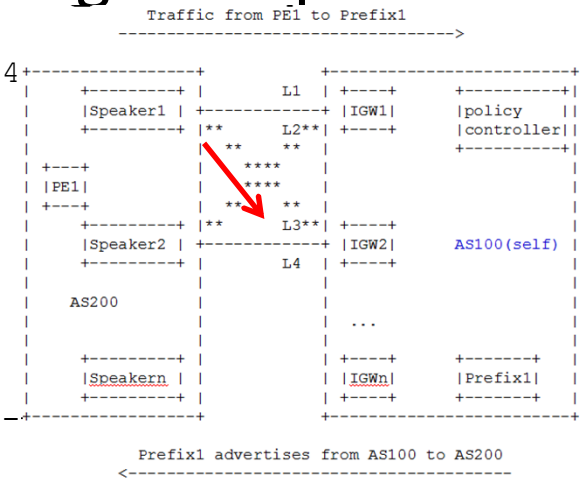

☐EBGP peering:
- •Speaker1---L1---IGW1
- •Speaker2---L2---IGW1
- •Speaker1---L3---IGW2
- •Speaker2---L4---IGW2

☐Requirement:
- •Administration only on AS100
- •Traffic enter AS100 through L3

- ☐ As required in the case, traffic from PE1 to Prefix1 need to enter through L3, so IGWs except IGW2 should prepend ASN list to Prefix1 when populating to AS100.
- ☐ As shown in left figure, community "PREPEND N TIMES TO AS" and "Exclude Target(s) TLV" are be used.

Application (2)

Outbound traffic control

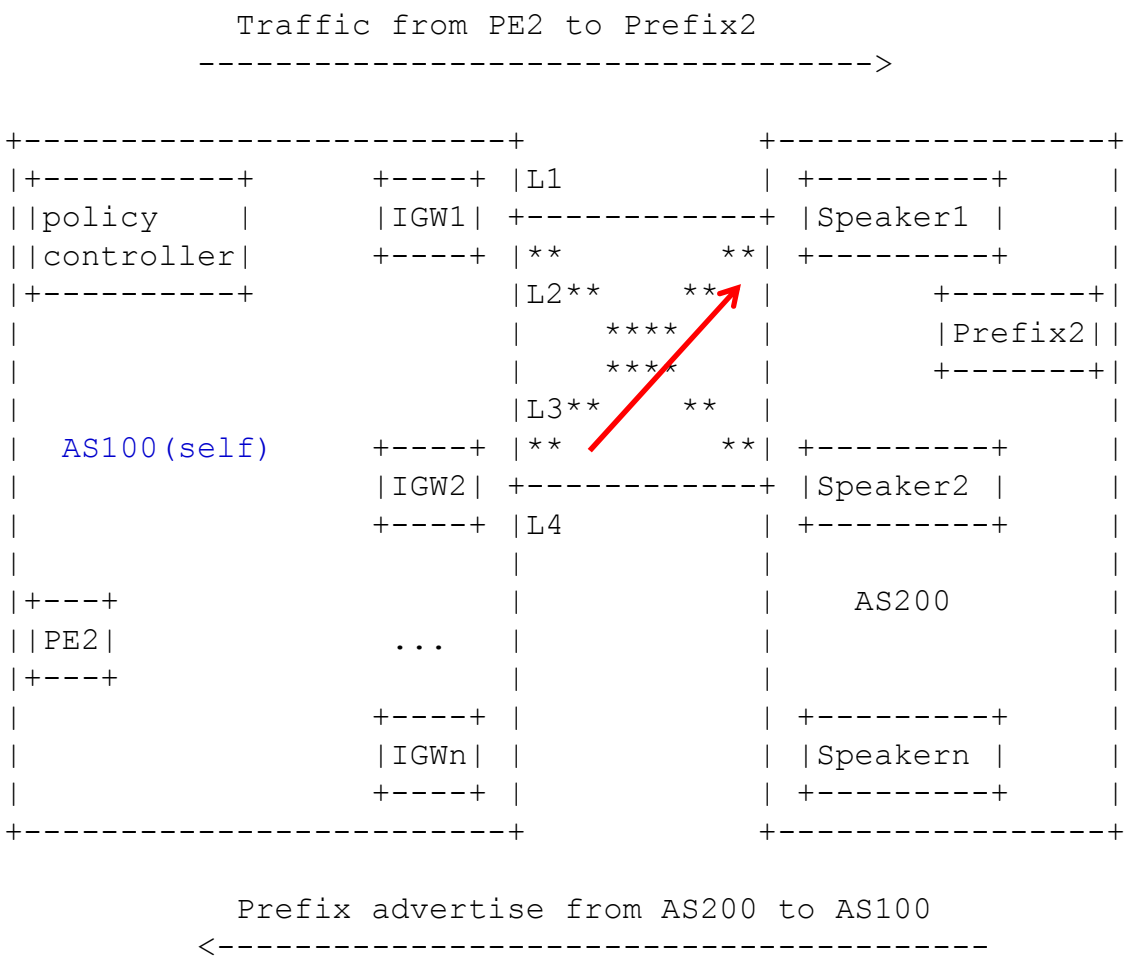

- EBGP peering:
  - IGW1---L1---Speaker1
  - IGW1---L2---Speaker2
  - IGW2---L3---Speaker1
  - IGW2---L4---Speaker2
- Requirement:
  - Administration only on AS100
  - Traffic exit through L3

Traffic Direction

Encoding Example (2)

Outbound Traffic Control encoding example

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Container Type 1 (1)                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|1 0 0 0 0 0 0 0|
+-+-+-+-+-+-+-+-+
| Hop Count: 0  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Length: 36                                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Community: LOCAL PREFERENCE                                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Own ASN                                                     100 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Context ASN#                                                100 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|TargetTLV(1)   | Length:                                    11 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IPv4Neig(TBD) | Length:                                     8 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Local Speaker                                           #IGW2 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Remote Speaker                                       #Speaker1 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Param TLV (3) | Length:                                     7 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Integer   (4) | Length:                                     4 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Increment #                                                 100 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

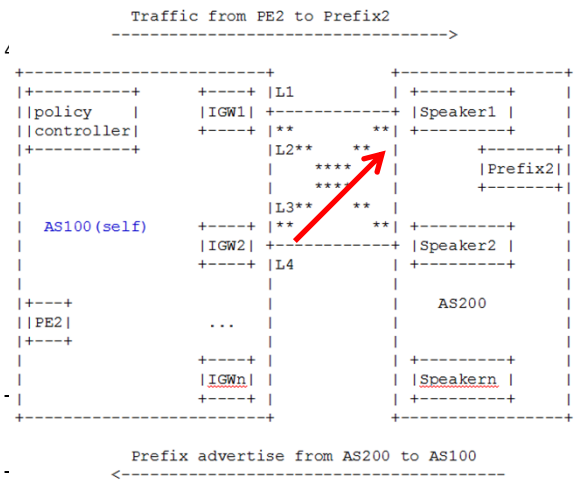

- EBGP peering:
  - IGW1---L1---Speaker1
  - IGW1---L2---Speaker2
  - IGW2---L3---Speaker1
  - IGW2---L4---Speaker2
- Requirement:
  - Administration only on AS100
  - Traffic exit through L3

- As required in the case, traffic from PE2 to Prefix2 need to exit through L3, so IGWs should prefer the route from IGW2 to Speaker1.
- As shown in left figure, community "LOCAL PREFERENCE" and "Target(s) TLV" are be used.

Next step

- Solicit comments on the alternative solutions.
- Refine this draft.
- Adding new use cases from operators.

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

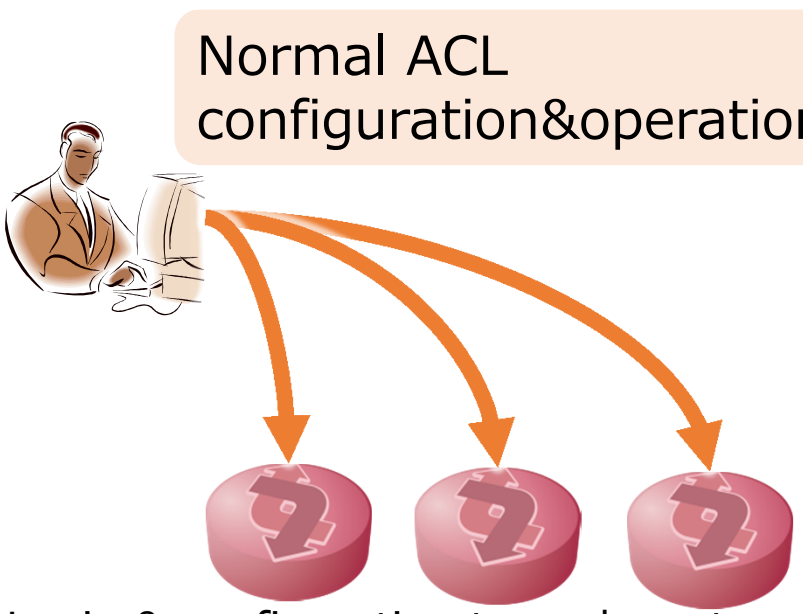

Login & configuration to each router
Too much work :(

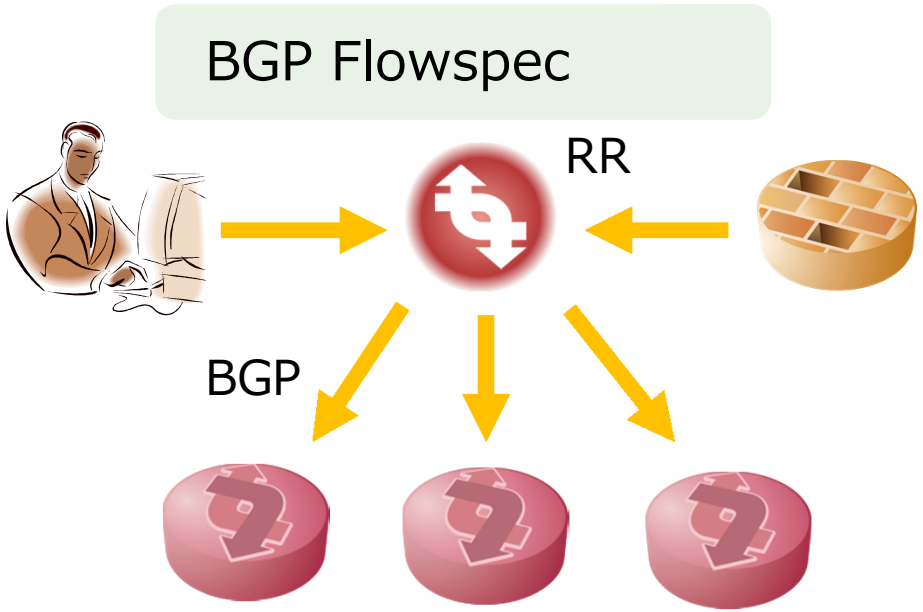

Easy to work together with security appliance

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

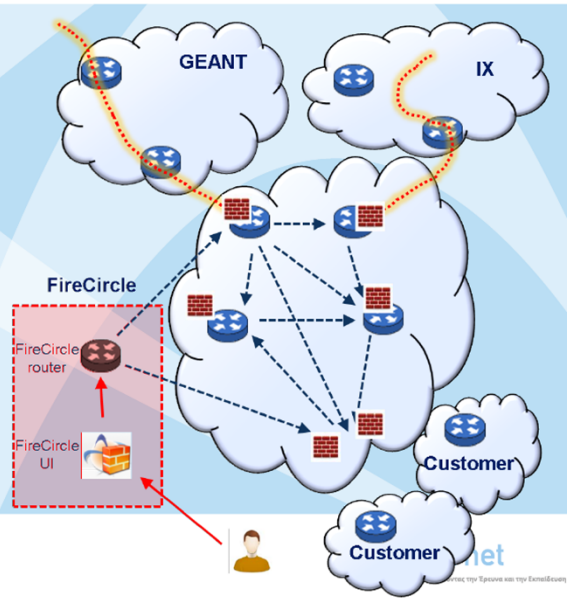

https://tnc2012.terena.org/core/presentation/41

NEO TELECOMS

Real life architecture

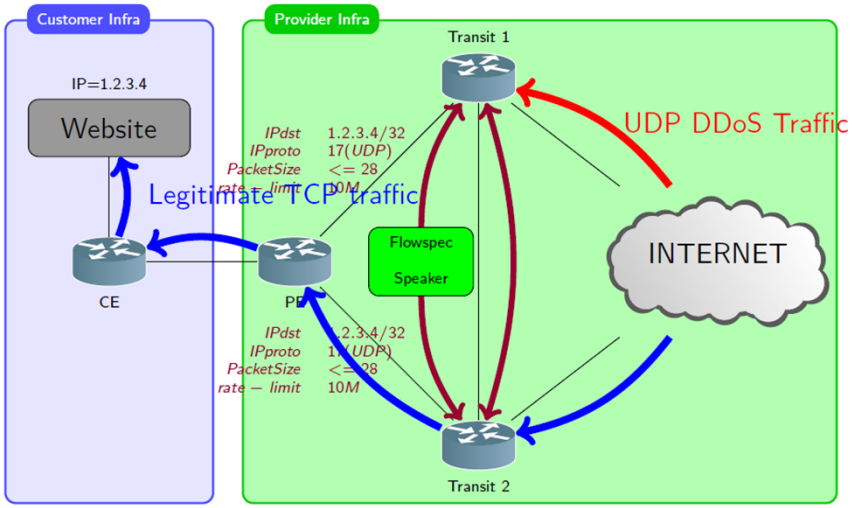

http://media.frnog.org/FRnOG_18/FRnOG_18-6.pdf

Use case

GRNET

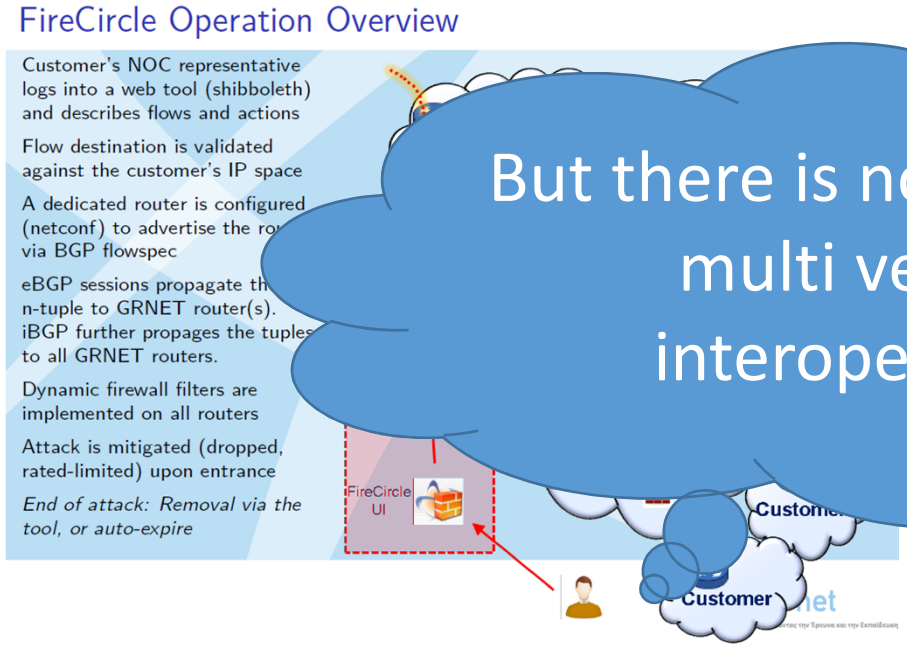

NEO TELECOM

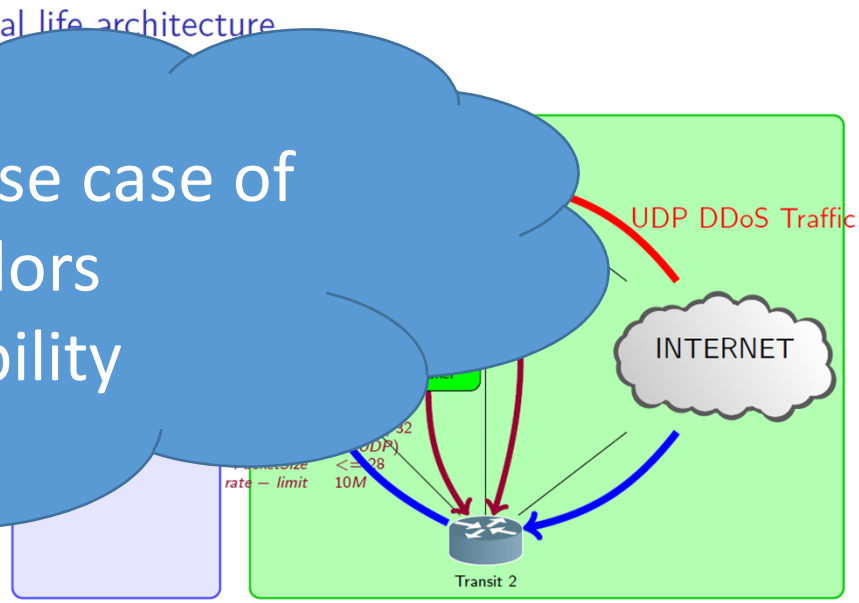

But there is no use case of multi vendors interoperability

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

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

Interoperability test topology @Interop Tokyo 2015 ShowNet

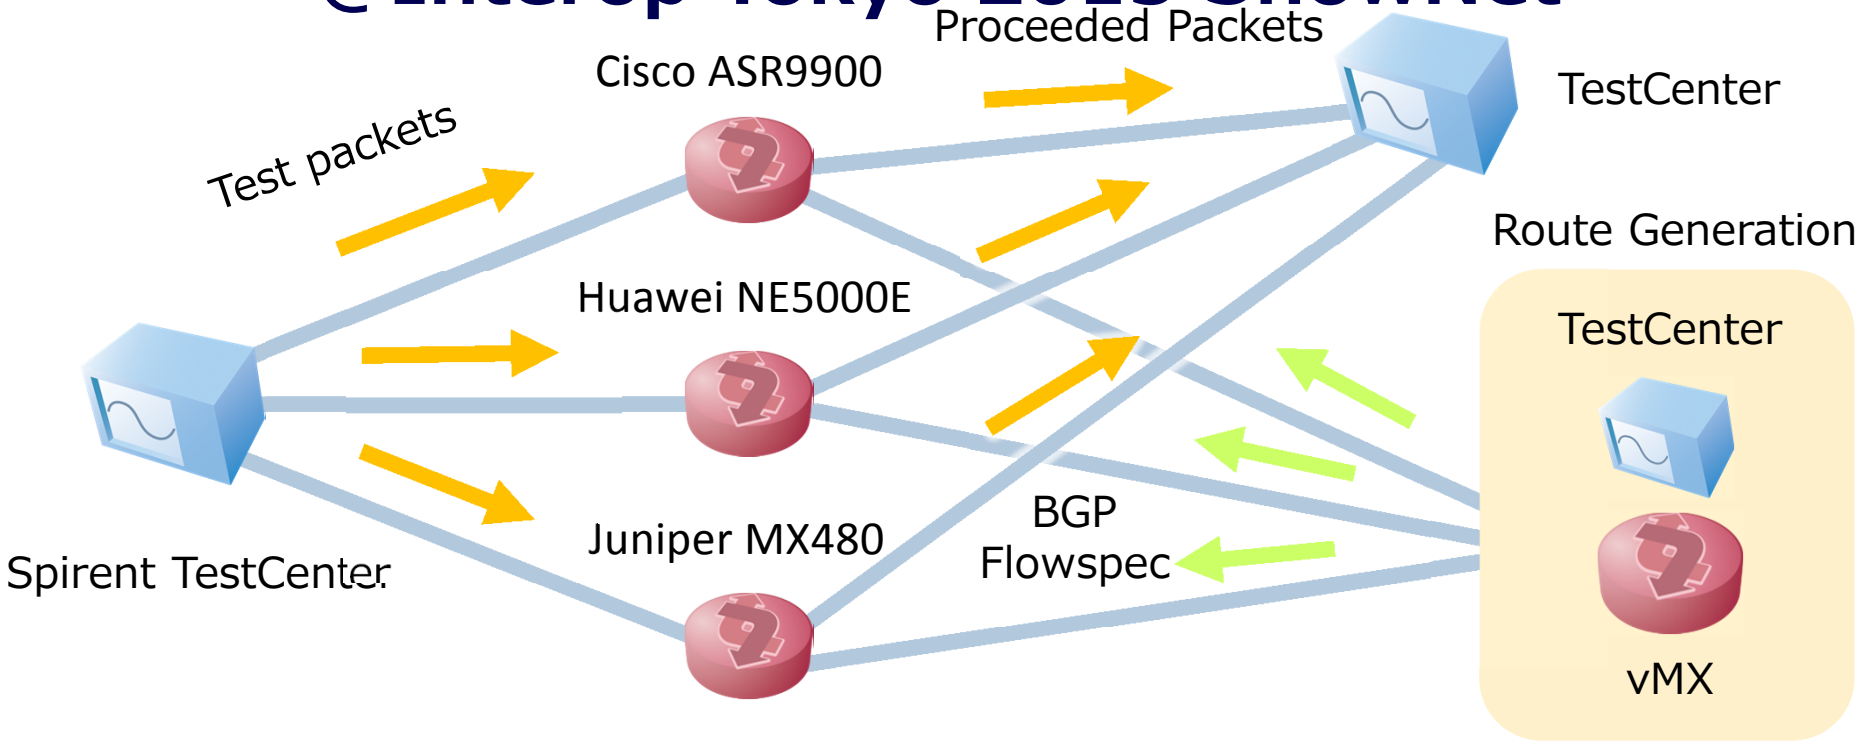

Test result BGP Flowspec Action rule

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

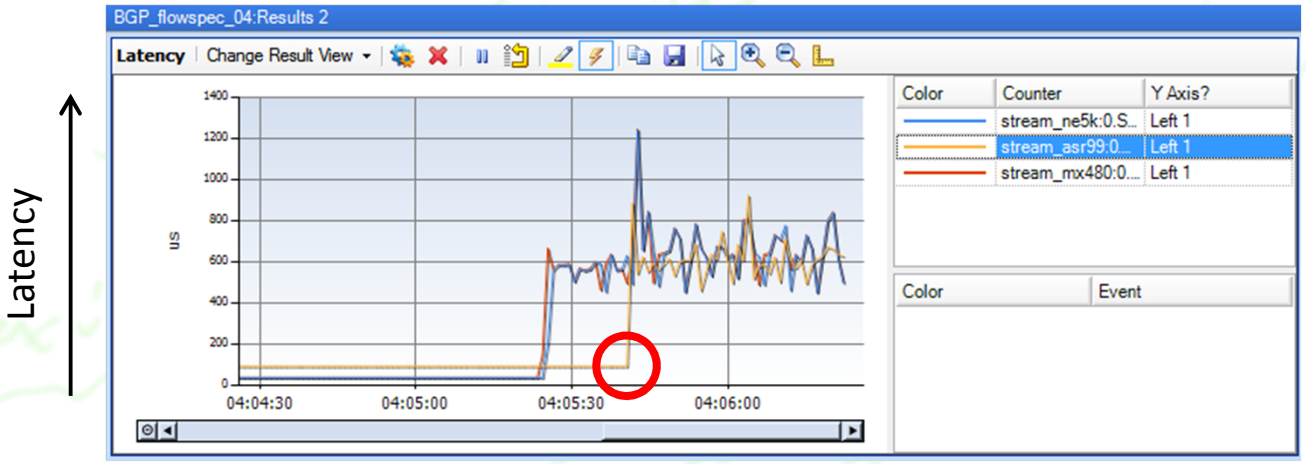

- Confirmed by measuring packets latency after redirecting (it's not caused by degradation of forwarding functionality of the router)
- ASR99xx took about 10 sec for processing after Redirection action rule injection. In case of withdrawn, the change was immediately reflected to the forwarding process.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

Rate-limit

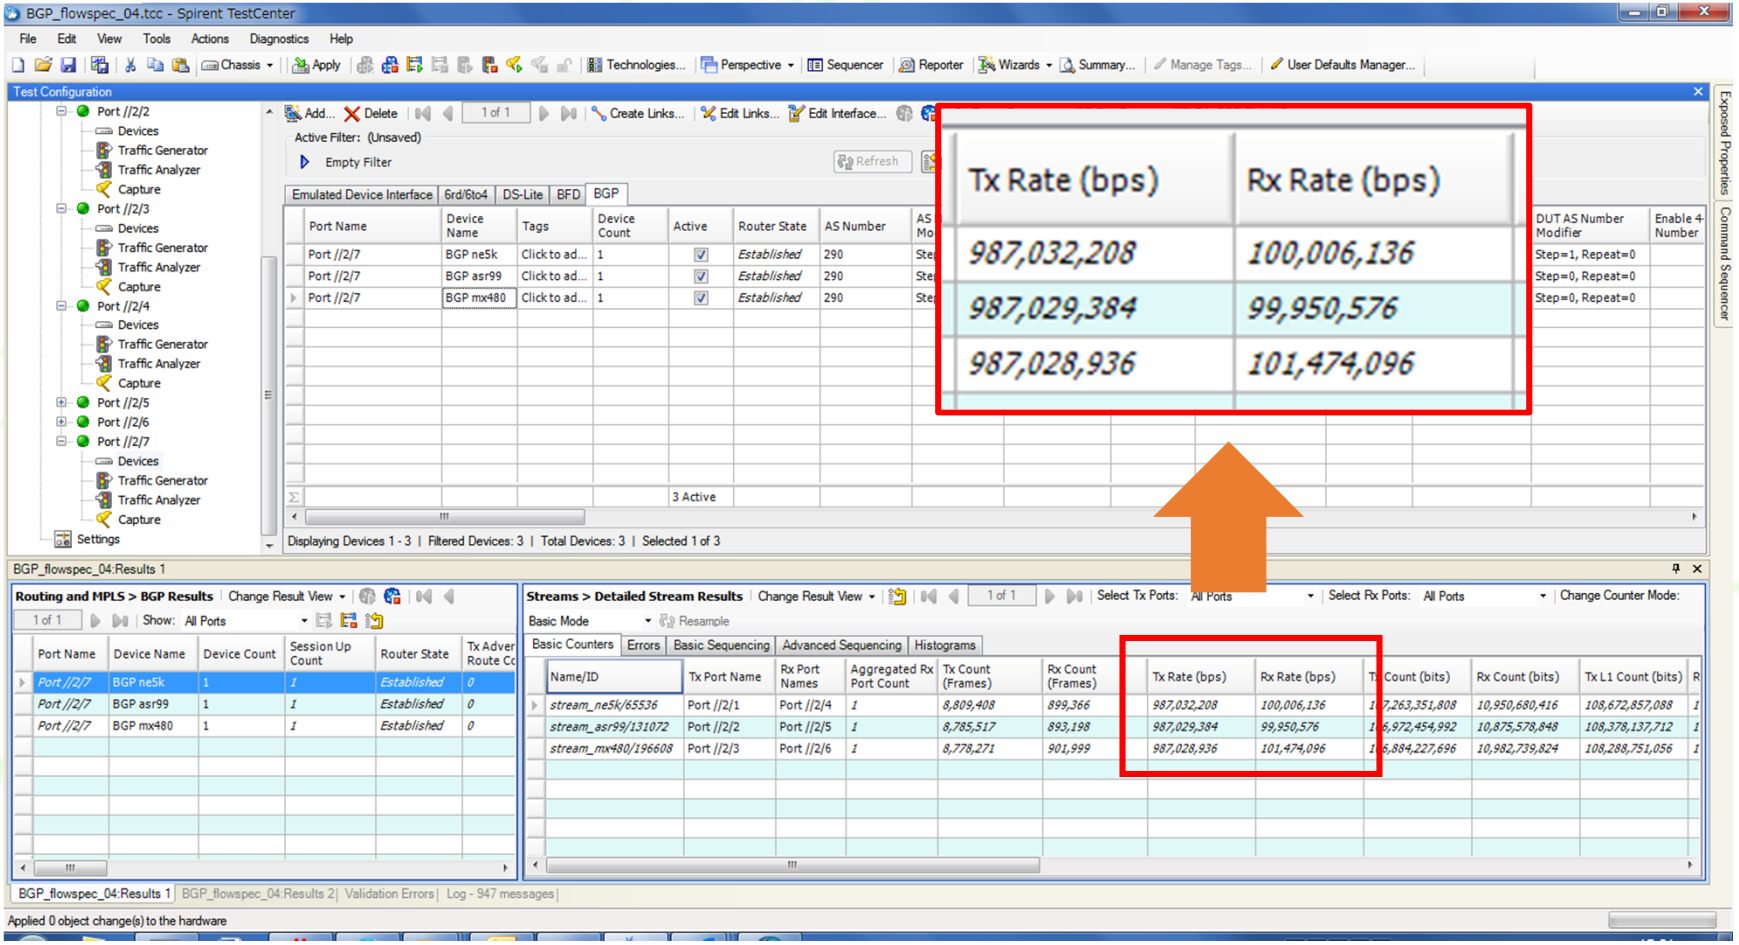
BGP_flowspec_04.tcc - Spirent TestCenter
Tx Rate (bps)
Rx Rate (bps)
987,032,208
100,006,136
987,029,384
99,950,576
987,028,936
101,474,096

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

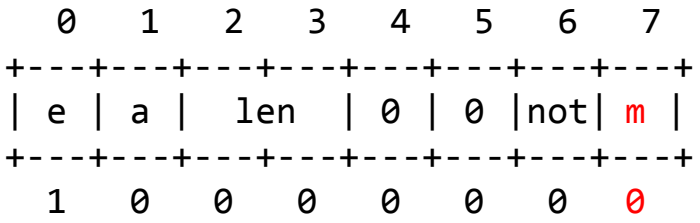

Cisco, Huawei

op=0x81

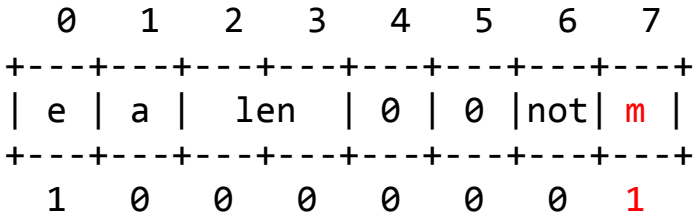

NE5000E treat as Invalid m=0

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

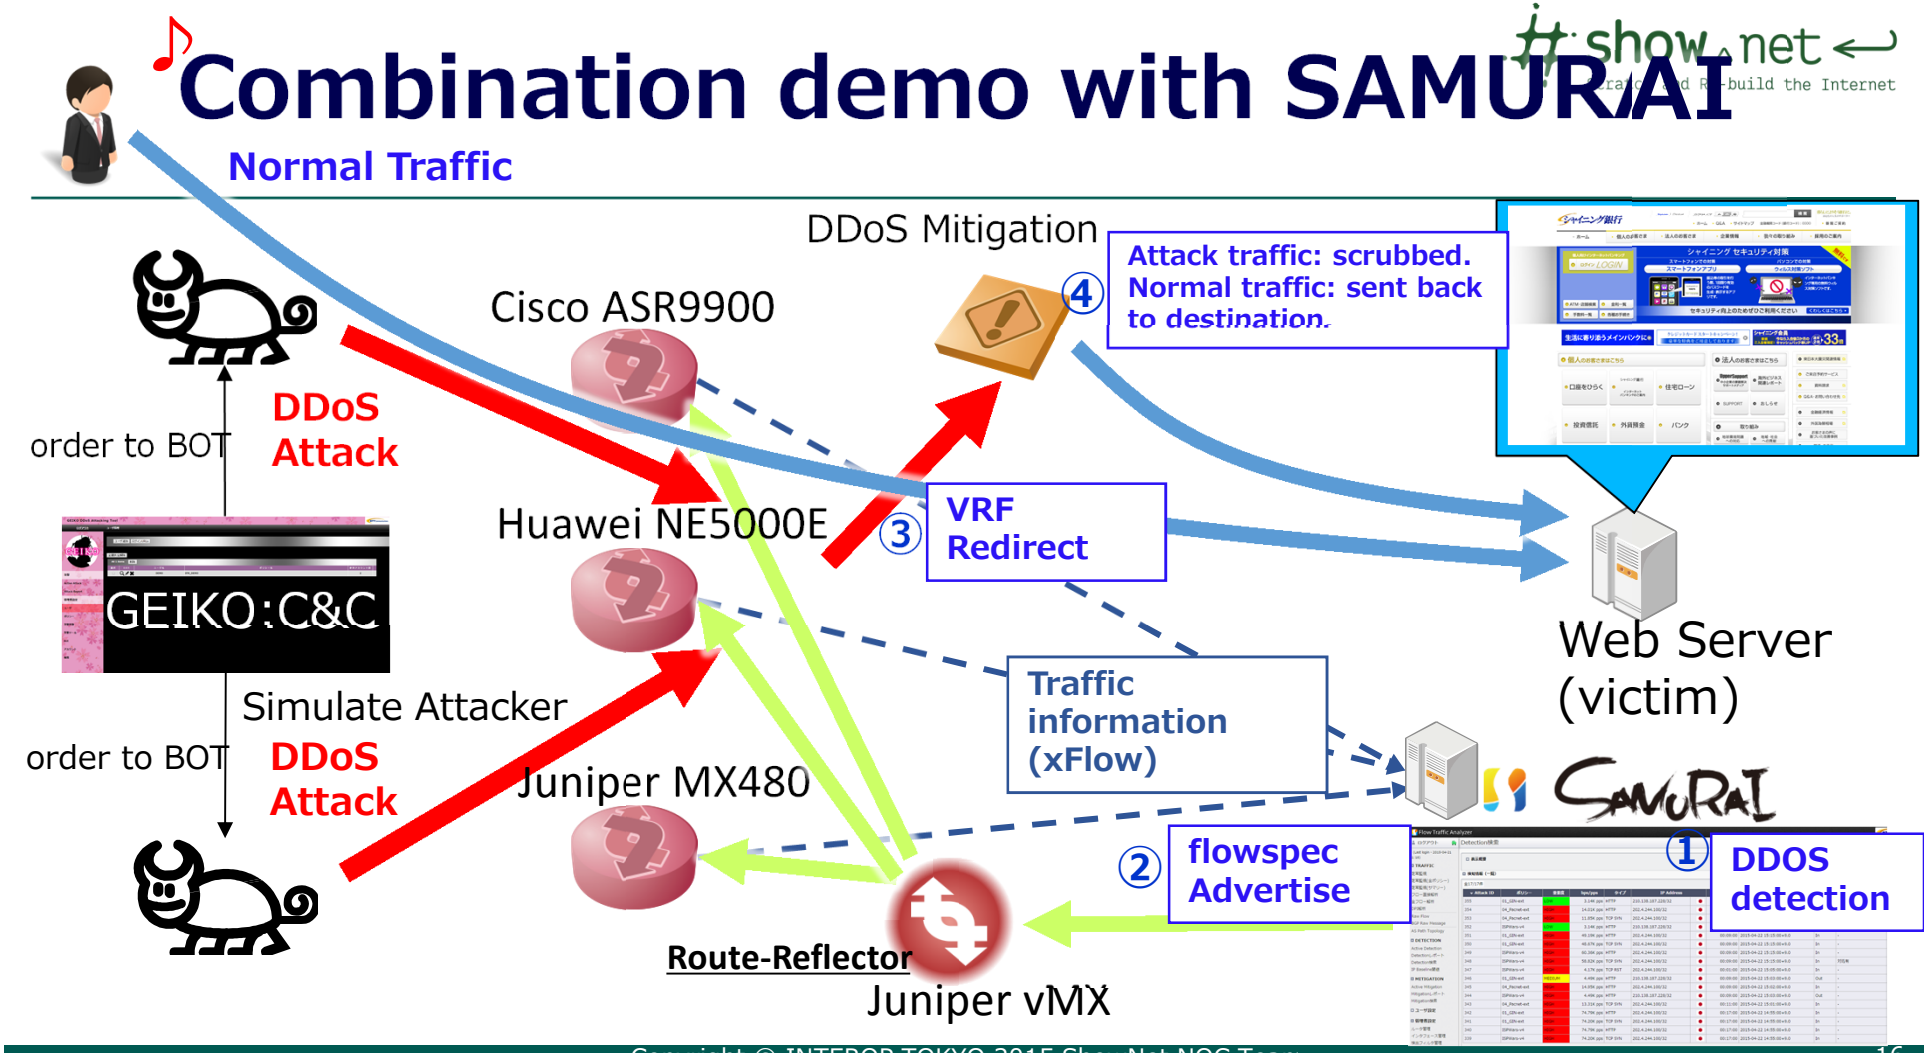

Summary

- 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

show net
Scratch and Re-build the Internet

Software Version

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