BGP Flow Specification – V2?
draft-hares-idr-flow-spec-combo-00.txt

Susan Hares

February 8, 2016

Agenda

- Overview
- BGP Flow Specification changes
- Precedence with Other Filters
- BGP Flow Specification Security
- BGP Flow Specification Yang module

Overview

- Review of RFC5575
- Why Precedence needed
- BGP Flow Specification: Ephemeral or Not?

Why Revise Flow Specification

- 2 IDR WG drafts + 9 proposals - Need rules for combination
  - draft-ietf-idr-flowspec-v6
  - draft-ietf-idr-flowspec-l2vpn
  - draft-eddy-idr-flowspec-packet-rate
  - draft-eddy-idr-flowspec-exp
  - draft-hao-idr-flowspec-nv03
  - draft-hao-flowspec-redirect-tunnel
  - draft-li-idr-flowspec-rpd
  - draft-liang-idr-bgp-flowspec-label
  - draft-liang-idr-flowspec-time
  - draft-litkowski-idr-flowspec-interfaceset
  - draft-vandevelde-idr-flowspec-path-redirect

- Should we create BGP Flow Specification V2?

Flow Spec (RFC5575) Review

- NLRI
  - For SAFI 133: IPv4 (AFI=1), IPv6 (AFI=2), L2VPN AFI=25)
  - For SAFI 134: IPv4 (AFI=1), IPv6 (AFI=2), L2VPN (AFI=25)
- Validation
  - Originator of flow spec = originator of best-match unicast route for destination embedded in NLRI,
  - No more specific unicast routes, when compared with Flow destination prefix, that have been received from different neighbor AS
- DoS and L2VPN doesn't fit

  ? – No destination check for SDN/VPN

  ? – does requirement this work for all Filters

  ? ROA's or BGP long-term

BGP Flow Specification is ECA Policy

- **ECA = Event – Match Condition - Action**
  - Flow-specification event = “packet reception”,
  - Condition – match filters in NLRI
  - Action – in Extended communities

- BGP Flow Specifications: last during BGP peer

Flow Specification Policy

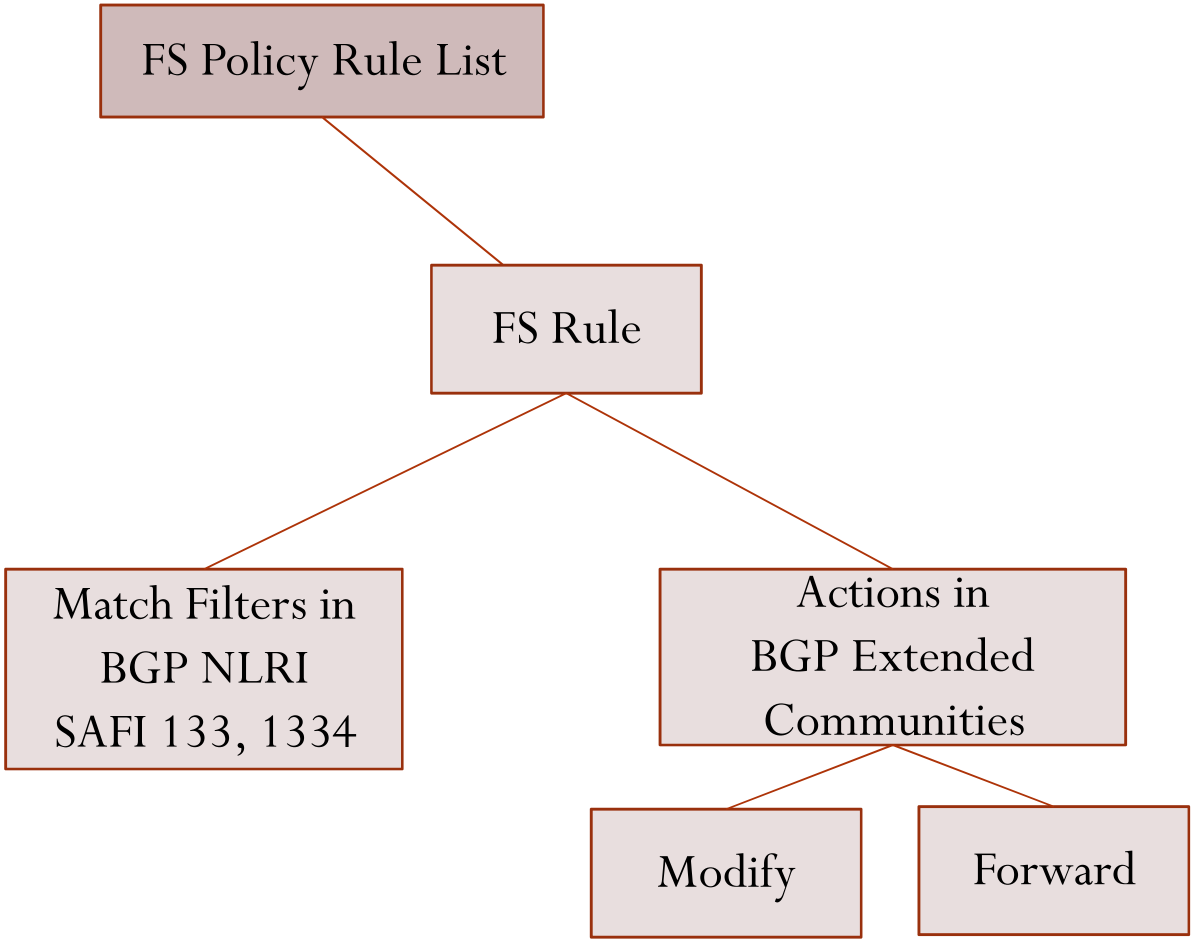
FS Policy Rule List
FS Rule
Match Filters in BGP NLRI SAFI 133, 1334
Actions in BGP Extended Communities
Modify
Forward

Why is precedence needed?

Precedence needed within BGP Flow Specification

- For filtering – Currently all
  - For ordering policies: use NLRI preference and administrative distance,
    - *Suggestion for change (Jeff Haas): Keep deployed FS, Updated Flow specification (address and rule order).*
  - For ordering filters – by Flow Specification type and precedence
- For action
  - No order currently, need to add order

BGP FS Filters types for RFC/WG documents

- RFC 5575 types/v6-draft
  1. Destination prefix
  2. Source prefix
  3. IPv4 protocol / IPv6 Next header
  4. Port (source or destination)
  5. Source port
  6. Destination port
  7. ICMP Type
  8. ICMP Code
  9. TCP Flags
  10. Packet length
  11. Traffic Class
  12. IPv4 Fragment
  13. IPv6 Flow ID

- L2VPN types
  14. Ethernet type
  15. Source MAC
  16. Destination MAC
  17. DSAP in LLC
  18. SSAP in LLC
  19. Control fields in LLC
  20. SNAP
  21. VLAN ID
  22. VLAN COS
  23. Inner VLAN ID
  24. Inner VLAN COS

BGP FS Proposed Filter types

- MF-1: NV03 Delimiter
  - Inner/outer header info
- MF-2: Virtual Network ID (VNID)
- MV-3: Flow ID (NVGRE Flow ID)

- MF-4 : MPLS LSP label or label stack
- MF-5: Interface Grouping
- MF-6: Time matches
- MF-7: Policy from IPv4 Neighbor
- MF-8: Policy from IPv6 Neighbor
- MF-9: Policy with AS Path

Are there others?

- Other types?
- Should we set a few types, and then create an Extended BGP Flow specifications
  - In another NLRI,
  - Or another BGP Attribute
  - (draft-li-flowspec-rpd)

BGP FS Filters: Precedence Rules (1)

Precedence logic for BGP Flow Specifications

(RFC5575, draft-idr-bgp-flowspec-l2vpn)

```
flow-rule-cmp (a,b)
{
  comp1 = next_component(a);
  comp2 = next_component(b);
  while (comp1 || comp2) {
   // component_type returns infinity on end of list
   if (component_type(comp1) &lt; component_type(comp2)) {
    return A_HAS_PRECEDENCE;
    }

   if (component_type(comp1) &gt; component_type(comp2)) {
    return B_HAS_PRECEDENCE;
   }
```

BGP FS Filters Precedence Rules (2)

```
// IP values)
 if (component_type(comp1) == IP_DESTINATION || IP_SOURCE) {
  common = MIN(prefix_length(comp1),prefix_length(comp2));
          cmp = prefix_compare (comp1,comp2,common);
          // not equal, lowest value has precedence
          // equal, longest match has precedence;
 } else if (component_type (comp1) == MAC_DESTINATION ||
     MAC_SOURCE) {
                    common = MIN(MAC_address_length(comp1),
                         MAC_address_length(comp2));
                    cmp = MAC_Address_compare(comp1,comp2,common);
                    //not equal, lowest value has precedence
                    //equal, longest match has precedence
          } else {
   common = MIN(component_length(comp1),
                         component_length(comp2));
           cmp = memcmp(data(comp1), data(comp2), common);
                    //not equal, lowest value has precedence
                    //equal, longest string has precedence
 }
 }
}
```

Flow Specification Actions

Approved Actions

(RFC 5575 & RFC 7674)

- Traffic rate in bytes (0x8006)
- Traffic Action (0x8007) with S(sample) T (terminal) flags
- Redirect to IP VPN via Route Target
  - RD 2 octet AS, 4 byte value (0x8008)
  - RD 4 octet IP, 2 byte value (0x8108),
  - RD 4 octet AS, 2 byte value (0x8208)

Proposed Actions

- (FA1) Traffic Rate in packets
- (FA2) Traffic Action with "R" for refer to more policy in BGP Attribute
- (FA3) Redirect to Tunnel
- (FA4) VLAN Action
- (FA5) TPID action
- (FA6) MPLS label action (push, pop, swap)
- (FA7) change validation to ROA or bgpsec-protocol
- (FA8a) interface set
- (FA8b) ACL+BGP FS

Default Precedence for BGP FS actions

- **Filters – AND**
  - 01-13: IP Protocol
  - 14-16: NVO3 matches [MF1-MF3]
  - 17: Segment ID
  - 18-29: MPLS [MF-4 + others]
  - 30-40: L2VPN matches (14-24)
  - 41: Interfaces matches (MF-5)
  - 42: Time matches (MF-6)
  - 43: IPv4 Neighbor
  - 44: IPv6 Neighbor
  - 45: AS Neighbor

Action

1. Alternate NLRI validation (FA-7)
2. Traffic rate in bytes (0x8006)
3. Traffic rate in packets (FA-1)
4. Traffic Action (0x8007)
5. Extended Traffic Action (FA-2)
6. Redirect to IP VPN (0x8008, 0x8108, 0x8208)
7. Redirect to tunnel (FA-3)
8. VLAN action (FM-4)
9. TPID action (FM-5)
10. Label Action (FM-6)
11. Interface Set (FM-8a)
12. Protocol Filter precedence (FM-8b)

Possible Conflicts

	Possible conflicts										
Action	Traffic rate Bytes	Traffic Rate Pkts	Traffic Action	Ext. Traffic Action	Redirect To IP VPN	Redirect to IP Tunnel	VLAN	TPID	Label	Intf Set	BGP valid
Traffic Rate Bytes		X									
Traffic rate Pkts	X										
Traffic action				X							
Ext. Traffic action			X								

Possible Conflicts

	Possible conflicts										
Action	Traffic rate Bytes	Traffic Rate Pkts	Traffic Action	Ext. Traffic Action	Redirect To IP VPN	Redirect to IP Tunnel	VLAN	TPID	Label	Intf Set	BGP valid
Redirect IP VPN						X	X	X	X	X	
Redirect Tunnel					X		X	X	X	X	
VLAN					X	X		X	X	X	
TPID					X	X	X		X	X	
Label					X	X	X	X		X	
Intf. Set					X	X	X	X	X		

BGP FS + Other Filter forwarding

- BGP is filter-Based forwarding
- Precedence between filter-forwarding

Packet/Frame Forwarding Filters

- **Where Forwarding Filters are created**
  - Configuration level: ACLs, PBRs
  - Box/module Ephemeral: I2RS
  - BGP Session Level: BGP Flow Specification

- **Filter-Based Forwarding is Minimalistic ECA Policy**
  - **Event** = packet reception on interfaces
  - **Match Condition** = Match on Filters
  - **Actions** – Modify packet, and Forward (or Drop)

- **Filters should have Yang data modules aligned**
- **Should this impact how BGP Flow filters are passed?**

Precedence between Flow Filters

- Why needed:
  - draft-litkowski-idr-flowspec-interfaceset proposes

    Really two actions

    - Apply policy to group of interfaces
    - Combine ACL + BGP Flow Specification filtering

  - Need Default Precedence + Policy Preference between:
    1) BGP Flow Specification (BGP Session Ephemeral)
    2) I2RS Filter Based RIB (Reboot Ephemeral)
    3) Filter-Based forwarding (aka Policy Routing) – configuration
    4) ACL – configuration

  - Propose Most dynamic (1st) to least dynamic (1-4 above)

BGP Security Upgrade for BGP FS

- BGP Flow Specification – pre-dates ROA and BGPSEC
- Validation using ROA
  - If have ROA: Use to validate transmitter of BGP FlowSpec along with Best-match unicast route for destination (IPv4 or IPv6)
    - ? Do we need no more specific route?
  - If no ROA: Best Match unicast route + no more specific routes

- Bgp-sec + ROA
  - Use signature to secure path of packet
  - Use ROA to secure ROA
  - Use Best match for unicast route for destination

Alternatives for BGP Flow-Spec V2

- Why change?
  - No order in BGP FS filters or BGP FS Action
  - Default precedence does not fit all cases
- Forms BGP Flow-Specification can take:
  - BGP attribute with ordered match filters and ordered actions: (draft-li-idr-flowspec-rpd)
  - NLRI + Wide communities
  - NLRI – with filters and actions

New filter match with order)

```
match type [bit 1 - deny/permit] 0-permit, 1 -deny

+------------------------+
| match type  (2 octets) |
+------------------------+
| number of sub-TLVS     |
|             (2 octets) |
+------------------------+
| sub-TLVs (variable)    |
| +====================+ |
| | order (2 octets)   | |
| +--------------------+ |
| | type (2 octets)    | |
| +--------------------+ |
| | length (2 octets)  | |
| +--------------------+ |
| | value (variable)   | |
| +====================+ |
+------------------------+
```

New Action with order

```
+--------------------------+
| Action order (2 octets)  |
+--------------------------+
| Action type (2 octets)   |
+--------------------------+
| Action length (2 octets) |
+--------------------------+
| Action Values (variable) |
+--------------------------+
```

Discussion Question 1

Should we go toward a BGP Flow Specification version 2 with ordered filters and ordered actions?

If so what format:

a) BGP Attribute

b) BGP NLRI with order + actions

c) BGP NLRI + BGP Wide Communities

d) Another Format?

Summary of January 11th discussion

- Why expand Flow Specification
  - Uses: DoS prevention, SDN/NFV, I2NSF
  - Need ordering for flow Specification
  - True Inter-Domain not as common within Provider with multiple AS-es
- If new mechanism, what about old?
  - Eventually Deprecate old, but allow side-by-side
  - Open Capability separate for New/Old

BGP Flow Specification Yang modules

- Yang Modules
  - BGP Flow specification similar to I2RSYang modules
  - Policy Based Routing (config)

BGP Flowspec vs. I2RS Filters

```
 Table 8 - comparison of Yang Data models
+-------------+----------------------+-----------------------+
| component    | BGP Flow Spec       | I2RS FB-RIB  +        |
|              | Yang                | Packet-ECA Yang       |
+==============+=====================+=======================+
|Policy        |flowspec-policy*     |group* [group-name]    |
| +-name       | [policy-name]       |                       |
| +-vrf        |+-rw vrf-name        | +-rw vrf-name         |
| +-AFI        |+-rw address-family  | +-rw address-famil    |
| +-rules      |+-rw flowspec-rule*  | +-rw group-rule-list  |
|              ||  [rule-name]       | | [rule-name]         |
|  +-rule-name ||+-rw rule-name      | |+-rw rule-name       |
|  +-rule-order||+-rw traffic-filters| |+-rw rule-order      |
|              ||+-rw traffic-actions| +-rw eca-rules        |
|              |                     | | [order-id rule-name]|
|              |                     | | +-rw installer      |
|              |                     | | +-rw eca-matches    |
|              |                     | | +-rw eca-qos-actions|
|              |                     | | +-rw eca-fwd-actions|
+--------------+---------------------+-----------------------+
```

Bgp Flow Spec vs I2RS Filters

```
+------------+----------------------+-------------------------+
| component  | BGP Flow Spec        | I2RS FB-RIB             |
|            | Yang                 | Packet-ECA Yang         |
+============+======================+=========================+
|opstate     |flowspec-state        |ietf-fb-ribs-oper-status |
| +-rib      |+-ro flowspec-rib     |+-ro fb-rib-oper-status* |
|            |  |                   |  +-ro fb-rib-name       |
|  +-groups  |  |                   |  +-ro group-status      |
|  +-rules   |  +-ro flowspec-entry*|  +-ro rules_opstate     |
|    [index] |    [index]           |  [rule-order, rule-name]|
|            |                      |  |                      |
```

Bgp Flow Spec vs I2RS Filters

```
+------------+----------------------+-------------------------+
| component  | BGP Flow Spec        | I2RS FB-RIB             |
|            | Yang                 | Packet-ECA Yang         |
+============+======================+=========================+
| +-rules    |+-ro flowspec-stats*  |  +-ro rules_opstats     |
|            |                      |  [rule-order, rule-name]|
|            | +-ro vrf-name        |                         |
|            | +-ro address-family  |                         |
|            | +-ro flowspec-rule-  |                         |
|            | |    stats           |                         |
|            | |                    |                         |
|            | |+-ro traffic-filters|                         |
|            | |+-ro traffic-action |                         |
|            | |+-ro classified-pkts|  | +--ro pkts-match     |
|            | |                    |  | +--ro pkts-modified  |
|            | |+-ro drop-pkts      |  | +--ro pkts-dropped   |
|            | |+-ro drop-bytes     |  | +--ro bytes-dropped  |
|            |                      |  | +--ro pkts-forwarded |
|            |                      |  | +--ro bytes-forwarded|
+------------+----------------------+-------------------------+
```

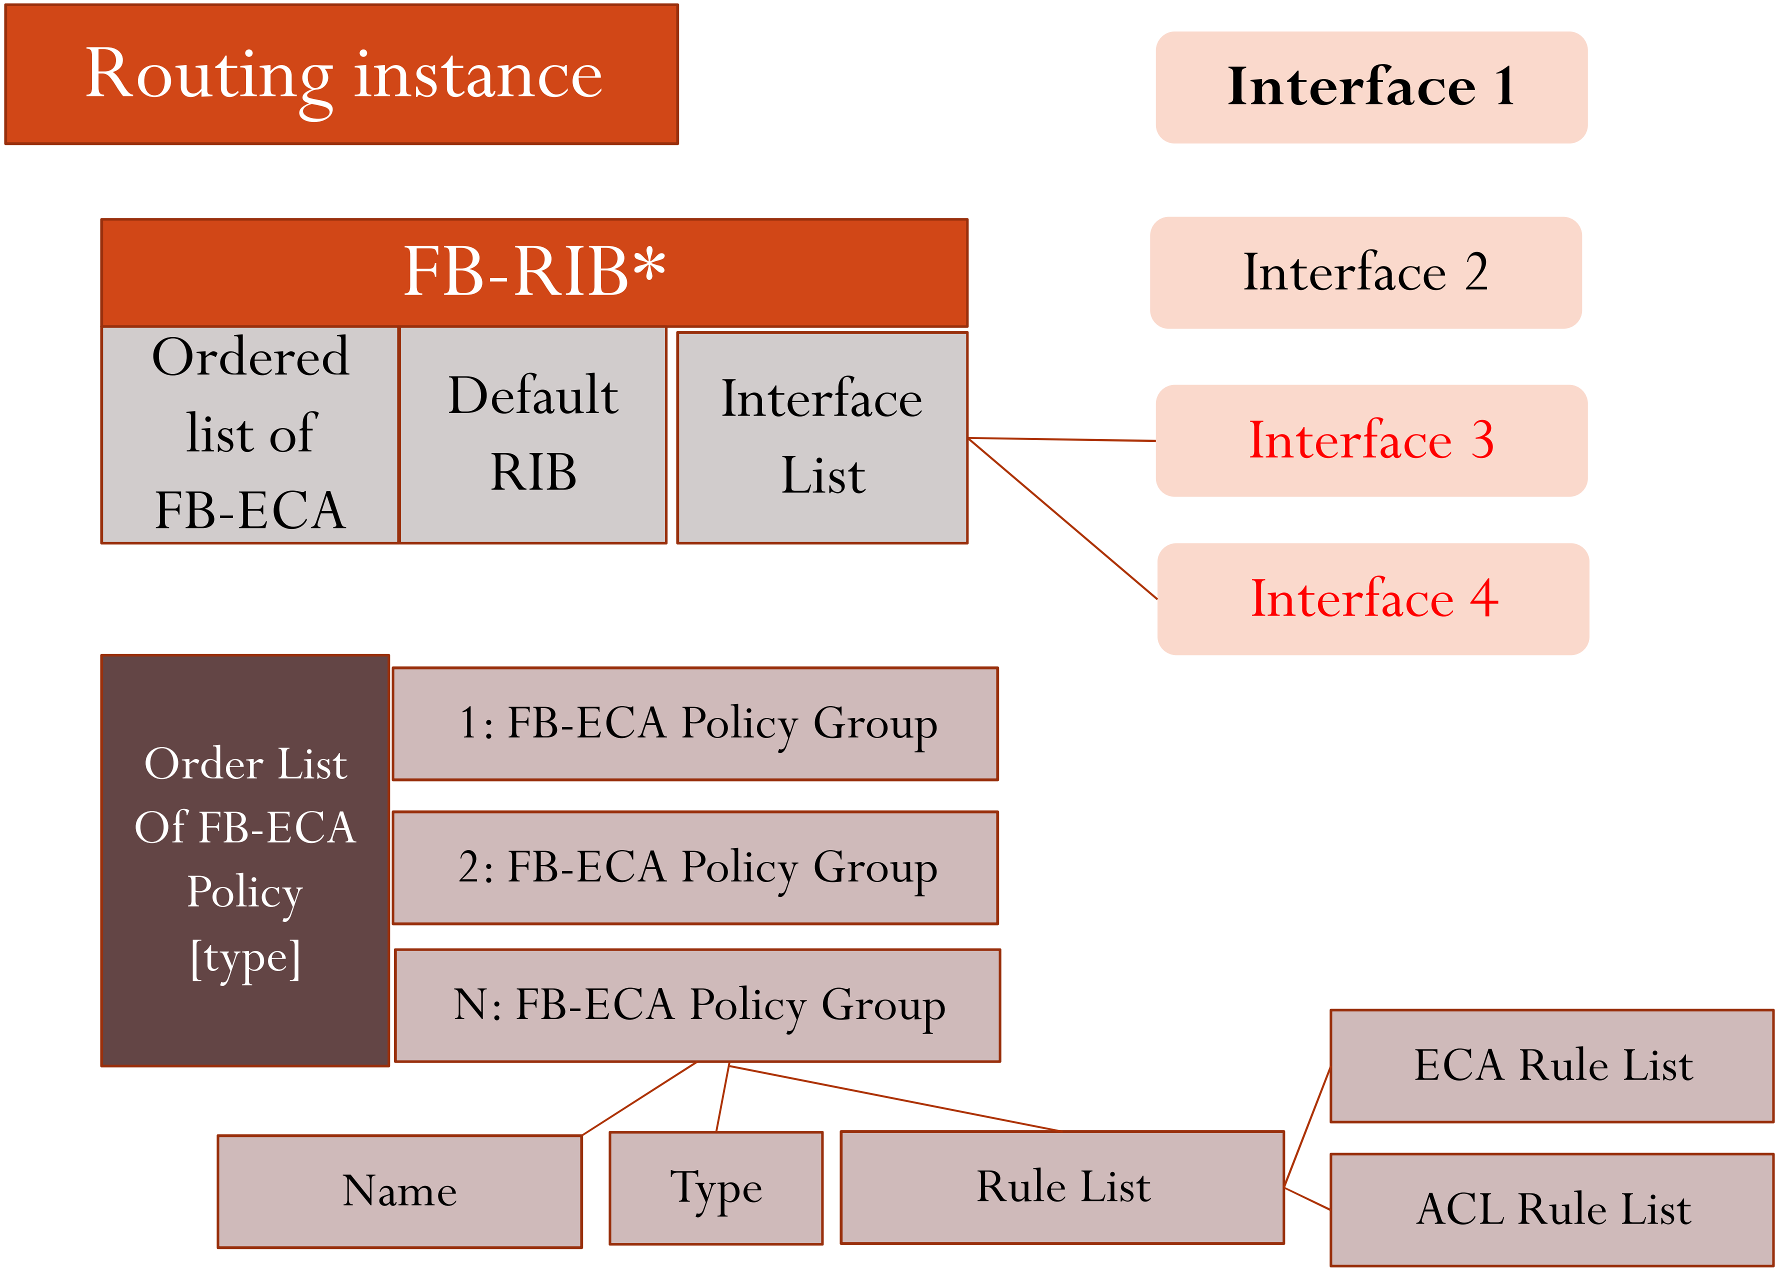

Routing instance

FB-RIB*

Ordered list of FB-ECA

Default RIB

Interface List

Interface 1

Interface 2

Interface 3

Interface 4

Order List Of FB-ECA Policy [type]

1: FB-ECA Policy Group

2: FB-ECA Policy Group

N: FB-ECA Policy Group

Name

Type

Rule List

ECA Rule List

ACL Rule List

FB-Rule List

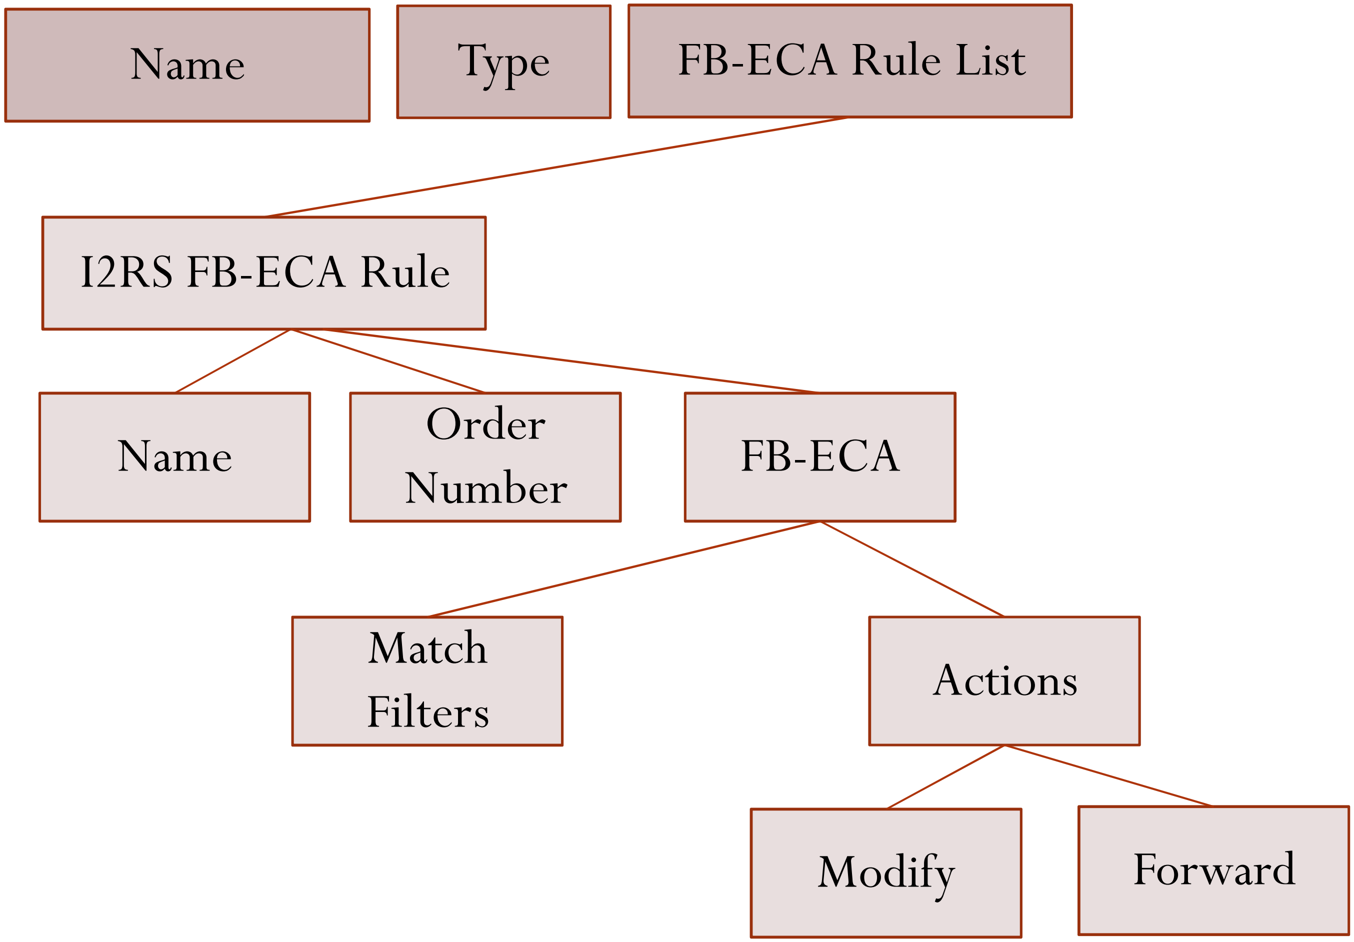

Name

Type

FB-ECA Rule List

I2RS FB-ECA Rule

Name

Order Number

FB-ECA

Match Filters

Actions

Modify

Forward

Filters in I2RS FB-RIB (hares-i2rs-pkt-eca-policy)

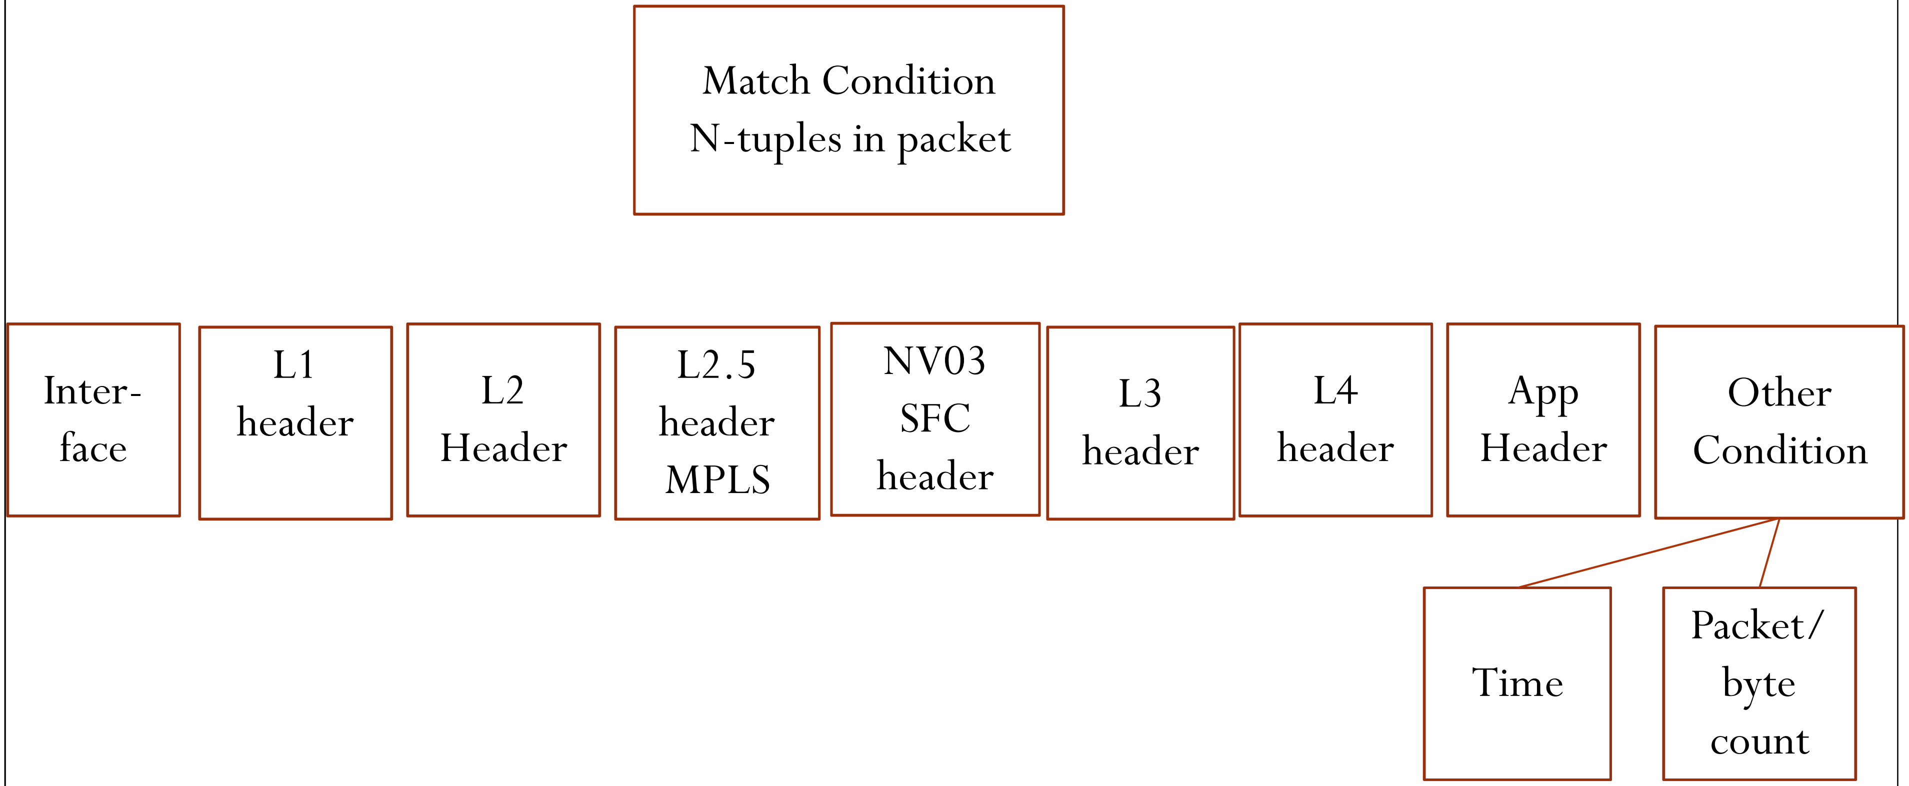

Discussion

Should we align all the Yang Modules for Filter-Based RIBs (config (aka policy routing), BGP, I2RS) ?

Backup slides

Discussion from 1/11/2016

- Should we have a successor to Flow-spec SAFI?
  - Action Criteria: IP Redirect (do-able) with 2 feature; Combination become with Actions is tricky;
    - Choice: combination
    - Precedence: better to specify, but will need to consider actions in combination
    - Redirect actions – interact with each; Modify actions interaction;
    - Traffic filters may
  - Match filters – as AND probably

Discussion Notes from 1/11/2016 (2)

- Flow Specification
  - Combination or separate Flow Spec
  - Rule ordering is reason for Flow-Spec 2,
    - Non-firewall, no SDN –may work
    - Firewall, SDN will not work without the ordering
  - Combination of the two flow-specification
    - If keep 2 SAFIs – two Flow-Specs into the future.
    - Ideal, v2 would have package with it – Date to deprecate V1 – real world doesn't probably won't allow it,
  - Agree with Jeff on backward compatibility
    - [wes] No way to tell which enhancement supported with out pre-knowledge,
    - [Jeff]: We do not have way to discover capabilities
    - [Robert]: We have this problem with the
    - [Jeff]: Redirect IP – possible that flow-specification – action (what does the implementation do with it).
  - Inter-domain flow-specification – not common
    - Service portals rather than inter-AS Flow specification
    - Redirect IP – within a single Provider – within a specific Provider

Discussion Notes from 1/11/2016 (3)

- Centralized mode –
  - Some flow specifications are only centralized controller and not distributed (Lucy Yong)
  - Some have two controllers (DDoS) and another (flow-filters)
    - Need to have precedence of the rules and then fall through (Jeff)
    - SDN (rule), and then the flow-specification rule
    - This requires a flow-specification v2 (jeff) because the existing things do not allow the flow-specification
    - Some the actions may only be appropriate to the list
  - Filter-based RIB
    - Precedence, fall-through – rule chains make sense
    - Take I2RS Filter-Based RIB

Discussion Notes from 1/11/2016 (4)

- Implementation of I2NSF
  - Controller tells the order of the rules
  - Can IDR provide this as well.
  - [Jeff]: More specific hosts, flow specification (longest prefix match will work)
  - [Linda]: Most specific

Discussion Notes from 1/11/2016 (5)

- John Schiel – flow spec rules that have precedence and ordering in flow specification rules.