FSv2 Choices

Susan Hares and Jeff Haas

draft-ietf-idr-fsv2-ip-basic-04

draft-hares-idr-fsv2-more-ip-filters-05

draft-hares-idr-fsv2-more-ip-actions-04

Topics

3 Questions to

- Minimal Compliant: FSv2 – AFI/SAFI + Extended Community

Next steps

- Interactions with Multiple FS Actions per filter rule
- BGP Community Container Attribute – FSv2 Container

Letting Filter Defaults – just work

- User Order [1-n] required by some applications
- Other Applications – just want defaults to work
  - Filter types (TLVs) are just groupings of Filters
- Two approaches for default (Components + filter TLV types)
  - Simple order (option 1)
  - Frame/Packet order (option 2)

FSv2 NLRI Format

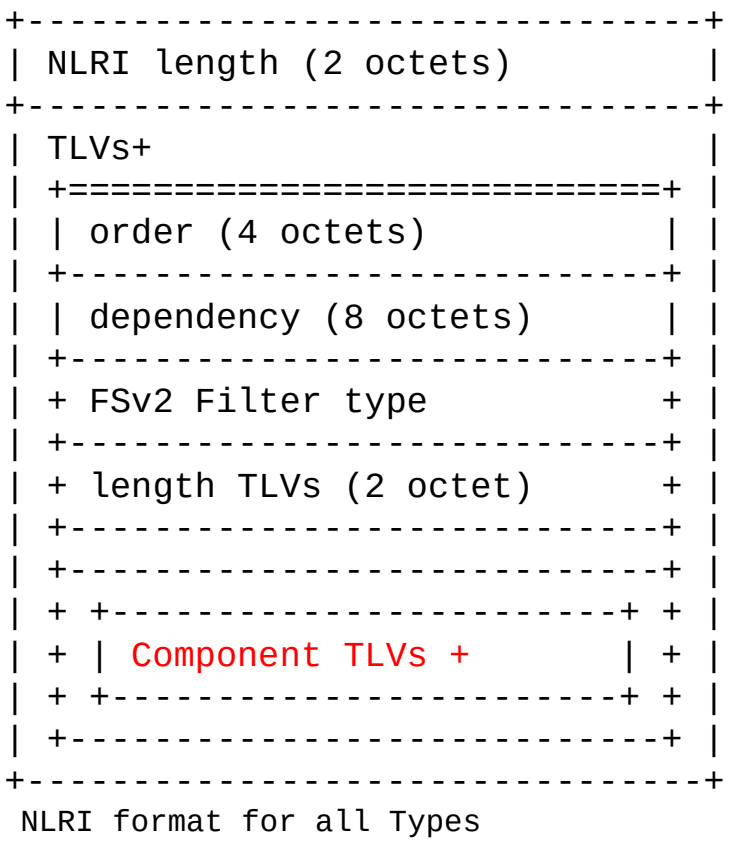

NLRI format for all Types

Component TLVs are:

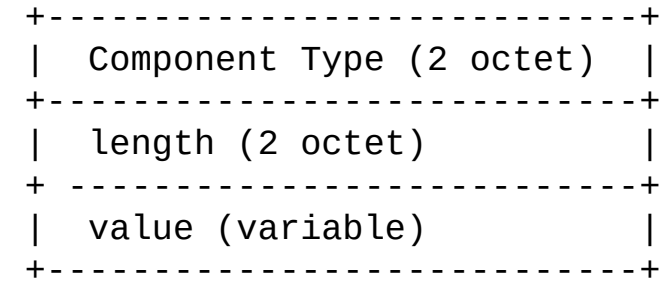

Question 1:
Filter type assignment – simple or frame order

Option 1 – simple ascending

0 - reserved

1 - IP Basic Filter rules

2 - IP Extended Filter Rules

3 - MPLS Filter Rules

4 - L2 traffic rules

5 - SFC Traffic rules

6 - Tunneled traffic

New Filter Types added at end

Option 2 – frame ordering

0 – reserved

50 - MPLS Traffic Rules

100 - L2 traffic rules

150 - SFC Traffic rules

200 - Tunneled traffic

256 - IP Basic Filter Rules

280 - IP Extended Filter Rules

New filters types added at frame location

Order: User order, then Filter types, then Components

Question 2:
Component Assignment - simple or packet order

Option 1

0 - reserved
1 - 13 - FSv1 components
14 - 63 - Extended Filter components
64 - 65 - MPLS Filter components
80 - 95 - L2 traffic filter components
96-120 - SFC Traffic components
121-140 - Tunneled traffic components
150+ - FCFS

Option 2 – frame ordering

0 reserved
1-13 FSv1
14- 49 Grouping before L2 (if, groups)
50- 99 L2 Traffic components
100-149 MPLS traffic components
150-199 SFC Traffic components
200-255 Tunneled traffic components
256-279 IP Basic Filter components
280-400 IP Extended Filter components

Consider where you would insert TTL for a packet

FSv2 Components At IETF-120 Option 1

FSv1 IP Component Numbers

1 - IP Destination prefix
2 - IP Source prefix
3 – IPv4 Protocol / IPv6 Upper Layer Protocol
4 – Port
5 – Destination Port
6 – Source Port
7 – ICMPv4 type / ICMPv6 type
8 – ICMPv4 code / ICPv6 code
9 – TCP Flags
10 – Packet length
11 – DSCP
12 – Fragment
13 – Flow Label

L2 Component Numbers [81-98]

1 - Ethernet type
2 - Source MAC
3 - Destination MAC
4 – DSAP in LLC
5 – SSAP in LLC
6 – control field in LLC
7 – SNAP
8 – VPAN ID
9 – VLAN PCP
10 – Inner VLAN ID
11 – Inner VLAN PCP
12 - VLAN DEI
13 - Inner VLAN DEI
14 – Src Mac Special bits
15 – Dst Mac Special bits
16 – RSN Mac Data unit
17 – Det. Latency Info

FSv2

14 – TTL
15 – SID in IPv6 Routing header
16 – NRP in Hop-by-Hop IPv6 header
17 – CAT ID (IPv6 header (?))
30 – Payload

Linked data Components (151- 180)

151 interface or interface group
152 Color
153 Time (or times)
154 AS or Set of Ases
155 Group and Sub-group

MPLS Component Numbers [1-2 or 64-65)

01 (64) MPLS Label Match-1 (label)
02 (65) MPLS Label Match-2 (Exp bits)

Tunnel Component Numbers [1-11 or 131-142)

01 – VN ID
04 – Cookie
05 – Tunnel header flags
06 – L2TP control version
07 – L2TPv3 Control Connection ID
08 – L2TPv3 Ns
09 – L2TPv3 Nr
10 – Protocol type
11 – GRE Sequence

FSv2 components At IETF-122 Packet order Option 2

IP Component Numbers (256-279)

256- Pre-TTL
257- IP Destination prefix
258- IP Source prefix
259- IPv4 Protocol / IPv6 Upper Layer Prot.
260- Port
261- Destination Port
262- Source Port
263- ICMPv4/v6 type
264- ICMPv4/v6 code
265- TCP Flags
266- Packet length
267- DSCP
268- Fragment
269-279: TBD

L3 Components for extended IP (280-400)

290– SID in IPv6 header
291- NRP in Hop by hop (IPv6)
295- CAT header

310 – Payload header

Packet Frame Filters

0-15 reserved for FSv1 16-31 – below L2
16 – Interface Group 18- AS Group
17 – time

L2 Component Numbers [50-99)

50 – L2 Filters

51- Ethernet type
52 - Source MAC
52 - Destination MAC
53 – 802.1Q tag
54 – Ethertype
60 – DSAP in LLC
61- SSAP in LLC
62- CTL in LLC
63- SNAP
71 – VLAN ID
72– VPAN ID
73- VLAN PC
80 – Inner VLAN ID
81 – Inner VLAN PCP
81- Inner VLAN DEI
82 – Src Mac Special bits
83 – DST Mac Special bits
84– RSN Mac data unit
85- Det. Latency info

MPLS Component Numbers [100-101]

101 - MPLS Label Match-1 (label)
102 - MPLS Label Match-2 (Exp bits)

Tunnel Component Numbers [200-250]

201 – VN ID
202 – Cookie
203 – Tunnel header flags
204 – L2TP control version
206- L2TPv3 Ns
207- L2TPv3 Nr
208 – Protocol type
209 – GRE

Sparse Components in Filter type TLV

- Option a: Type TLV

  If adding new components, a Type TLV groups a set of known components

  Type TLV = FSv1 (256), components allowed (257-268)
  Type TLV = Ext-FS (280), components allowed (256-268, 290-291, 310)

- Option b: Component that identifies filters supported

  Component assigned to identify filters sent

Option a + b – can co-exist

FSv2 Component For supported Filters (bit mask)

IP Component Numbers (257-28)

256- Pre-TTL
257- IP Destination prefix
258- IP Source prefix
259- IPv4 Protocol / IPv6 Upper Layer Prot.
260- Port
261- Destination Port
262- Source Port
263- ICMPv4/v6 type
264- ICMPv4/v6 code
265- TCP Flags
266- Packet length
267- DSCP
268- Fragment
269-279: Resrved

L3 Components for extended IP (280-400)

280- Filters

281 – Group (for L3)
290– SID in IPv6 header
291- NRP in Hop by hop (IPv6)
295- CAT header

310 – Payload header

Packet Frame Filters

0- FSv1 filters supported

1-14 Reserved for FSv1 15 – time
15 – group (interface) 16 – time

L2 Component Numbers [50-99)

50 – L2 Filters

51- Ethernet type
52 - Source MAC
52 - Destination MAC
53 – 802.1Q tag
54 – Ethertype
60 – DSAP in LLC
61- SSAP in LLC
62- CTL in LLC
63- SNAP
71 – VLAN ID
72– VPAN ID
73- VLAN PC
80 – Inner VLAN ID
81 – Inner VLAN PCP
81- Inner VLAN DEI
82 – Src Mac Special bits
83 – DST Mac Special bits
84– RSN Mac data unit
85- Det. Latency info

MPLS Component Numbers [100-101]

100- MPLS Filters

101 - MPLS Label Match-1 (label)
102 - MPLS Label Match-2 (Exp bits)

Tunnel Component Numbers [200-250]

200 – Tunnel Filters

201 – VN ID
202 – Cookie
203 – Tunnel header flags
204 – L2TP control version
206- L2TPv3 Ns
207- L2TPv3 Nr
208 – Protocol type
209 – GRE

FS Actions Plan

1. Actions - Use Extended Community Actions for FSv1 and FSv2

For New actions –just describe

- Describe which actions interact with
- Describe the order of actions
- (no protocol enforcement, vendor must test interoperability)

2. Work on prototypes for ordered actions

- Work with early prototypes + hackathons
- Quick feedback to authors + deployers

Minimal FSv2

Minimal FSv2 Compliant:

- 1 AFI/SAFI NLRI pair + FSv2 Extended Communities
- One filter group restricted to: only FSv1 filters in FSv2

Minimal FSv2 SRv6 compliance

- 1 AFI/SAFI NLRI pair + FSv2 Extended Communities
- One filter group with FSv2 filters needed

Background slides

FS Actions

- Existing Actions in Extended Community (RFC + WG)
- Analyzing Actions – categories, support, optionality
- Encoding actions

Actions in Flowspec RFCs + WG drafts

- **Defined in RFC 8955, Section 7**:

0x8006	traffic-rate-bytes (Section 7.1)	2-octet AS, 4-octet float
0x800c	traffic-rate-packets (Section 7.2)	2-octet AS, 4-octet float
0x8007	traffic-action (Section 7.3)	bitmask
0x8008	rt-redirect AS-2octet (Section 7.4)	2-octet AS, 4-octet value
0x8108	rt-redirect IPv4 (Section 7.4)	4-octet IPv4 address, 2-octet value
0x8208	rt-redirect AS-4octet (Section 7.4)	4-octet AS, 2-octet value
0x8009	traffic-marking (Section 7.5)	DSCP value

Redirect defined in RFC

- Redirect to an SFC path - RFC 9015

Redirect defined in WG drafts

- redirect-ip - draft-ietf-idr-flowspec-redirect-ip
- redirect to indirection ID – draft-ietf-idr-flowspec-path-redirect (IPv4)
  - draft-ietf0-idr-flowspec-path-redirect (IPv6)

Match Criteria Carried in Extended Communities

- draft-ietf-idr-flowspec-interfaceset permits interfaces to be matched by a group ID on a directional basis.
  - Discussion among the authors suggests that directionality might not be able to be enforced and we should drop it as a criteria.

- This isn't *quite* an Action, but perhaps it may be a necessary component for thinking about conflicts.

Categorization

- When trying to work address how conflicts should be addressed in FSv2, and how prior FSv1 implementations dealt with the conflicts, it was helpful to split the actions into categories of functionality.
- In general, an action in one category doesn’t conflict with actions in another category.
- Within the same category, conflict may occur.

Categories of Existing FSv1 Extended Communities

1. Limits on filters or actions (only on a set of interfaces)
   - Note - this may not be appropriate for ”actions” and needs discussion.
2. Rate limiting (bytes or packets)
3. Set DSCP value in IP Packet
4. Sample Packet (TAIS, sample)
5. Redirect traffic flow (to: VRF, VPN, SFC Path, redirect-id)
6. Stop Processing (TAIS, terminal)

Support

- Even when an implementation understands what an action is, it may not have support for it.
- Perhaps this lack of support is via configuration!

Optionality

- When should the lack of support for a given action be considered a problem?
  - Some actions may be *optional*. For example, setting DSCP as a “nice to have” to shape traffic in a network may not be a required thing. Thus, failure to be able to implement it might be okay.
  - However, perhaps DSCP is required for class-based-forwarding and thus for this specific FS route, it is *required*.
- The operator is the one that will know what is necessary for their FS routes.

Encoding

- The necessary inputs are:
  - What action types are to be implemented?
  - Is a given action optional or required?
  - Are there conflicts?
  - What do we do on failure?
- These can be encoded as *action tuples*. This is a list of actions that should have no conflicts where all required actions MUST be able to be installed for the FS route.

Interactions Between Actions (FSv2-EC)

- If Single actions per filter match, no problem.
- If multiple actions per filter match, actions can interact
  if actions **succeed** or **failure**

For Success: 3 simple steps to handle interactions of FSv2-EC

1. Analyze which category
2. Determine if one action per category
3. If more than one per category, do not install multiple actions

For Failure: if Best-effort, install.

If "all or none", skip unless can handle

Simple test case for Multiple Actions (Success + failure: Best Effort)

- Case 1:
  - Set DSCP value in packet header
  - Sample the traffic
  - Redirect to VRF
- Case 2:
  - Redirect to VRF
  - TAIS terminates filter processing
- Case 3
  - Redirect to VRF
  - Redirect to SFC

Case 1: No interactions

Case 2: No interactions

Case 3: Interactions between Actions, don't install

Simple test case for Multiple Actions (Success and all or none)

- Case 1:
  - Set DSCP value in packet header
  - Sample the traffic
  - Redirect to VRF
- Case 2:
  - Redirect to VRF
  - TAIS terminates filter processing
- Case 3
  - Redirect to VRF
  - Redirect to SFC

Case 1: No interactions

if DSCP can fail, don't install

Case 2: No interactions

Case 3: Interactions between Actions, don't install

Path Attribute for Actions

- List of:
  - Priority
  - Action Tuple:
    - Optional/Required
    - Action
    - Action Parameters

- In priority order, try an action tuple and see if all required actions can be implemented. If not, iterate.
- If no action tuples can be implemented, take the desired default action.

Possible TLVs in BGP FSV2 Container
(draft-hares-idr-fsv2-more-ip-actions)

1. Action Chain ordering – Multiple Actions (Success/Failure)
2. Interface set – operation actions or filters on a set of interfaces
3. Groups – set some grouping (filters or actions)
4. Traffic rate limits - parameter: bytes or packet)
5. Traffic Sample or Copy
6. Traffic mark
7. Redirect with sub-TLVs for Redirect

Next Steps

- Get Feedback
- Spin new versions of FSv2-IP-Basic + FSv2-More-IP Filters
- WG adoption + WG LC for FSv2 Extended Communities
- WG adoption call for FSv2-More-IP Filters
  and individual drafts
- Work on BGP Community Container for FSv2

Backup Slides

BGP Community Container attribute draft-ietf-idr-wide-bgp-communities-12

- User Order of Actions requires more bytes than Ext. Community
- Common Header for Attribute (Transitive/Confederation)