IDR FSv2 Interim Design Team 3+ 4

5/6/2024

Time: 10-11:30am

Agenda for Design Team 3
FSv2 Actions for Basic IP

---

0) Agenda Bashing + Introductions [5 minutes]

1) Short overview of FSv2 for Basic IP [5 minutes]

draft-hares-idr-fsv2-ip-basic (Sue Hares]

2) FSv2 Actions Formats – Ext Community + Community Path Attribute

draft-hares-idr-more-ip-actions (Sue Hares) [15 minutes]

3) FSv2 Action Discussion on Extended Communities [10 minutes]

Ordering + Conflicts

4) FSv2 User Ordering in Community Path Attribute [10 minutes]

Ordering + Conflicts + Dependency

Agenda for Design Team 4
FSv2 for Basic IP (30 minutes)

Existing Non-IP Filtering + Actions

1. MPLS + SFC – Sue Hares [5 minutes]

draft-ietf-idr-flowspec-v2, RFC9015

2. FSv1 L2VPN– Donald Eastlake [15 minutes]

draft-ietf-idr-flowspec-l2vpn

3 FSv1 Tunnels – Donald Eastlke [15 minutes]

draft-ietf-idr-flowspec-nv03

FSv2 – IDR Interims operating as Open Parallel design teams to break FSv2 into “chunks”

Team 1: Basic IP FSv2: Current IPv4/v6 filters + current actions + order

Team 2: More IP Filters FSv2 - Defining more IP filters to add to IP

Team 3: More IP actions FSv2 – Defining more actions + action sequences to add to

Team 4: Non-IP filters and actions FSv2

Just ask to join a team!

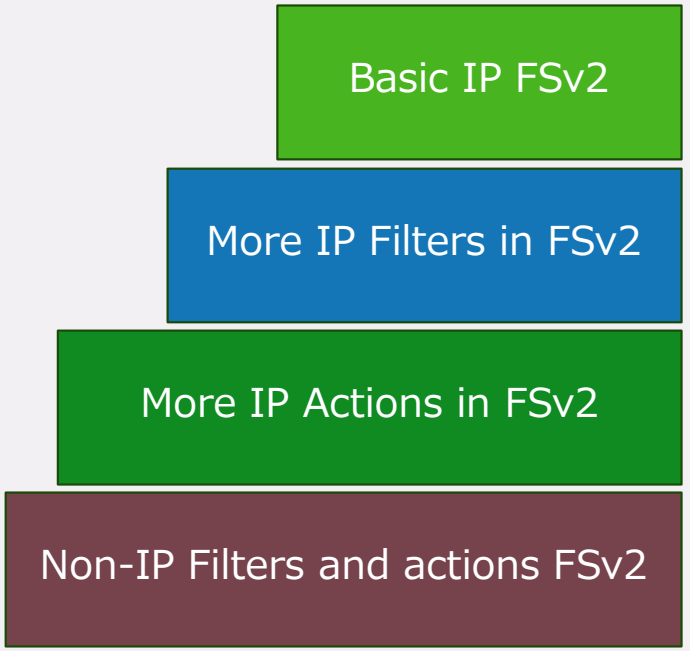

Interims for FSv2 Open Design Teams

Phase 1 – May - Start our work

- 4/29 – Design Team 1
- 5/06 – Design Team 1 + 2
- 5/13 – Design 3+ 4 – IP Action Design + Non-IP Filters/Actions

Phase 2 – June – Getting ready for IETF-120

- 6/03 – Design Team 1 – Review of FSv2 Basic IP
- 6/10 – Design Team 1 + 2 – Review of FSv2 - More IP Filters
- 6/17 – Design Team 3 + 4 – Review of IP Actions

IDR FSv2 Interim Design Team 3

5/13/2024

Time: 10-11:30am

FSv2 for Basic IP Review

draft-hares-idr-fsv2-ip-basic-01

Sue Hares

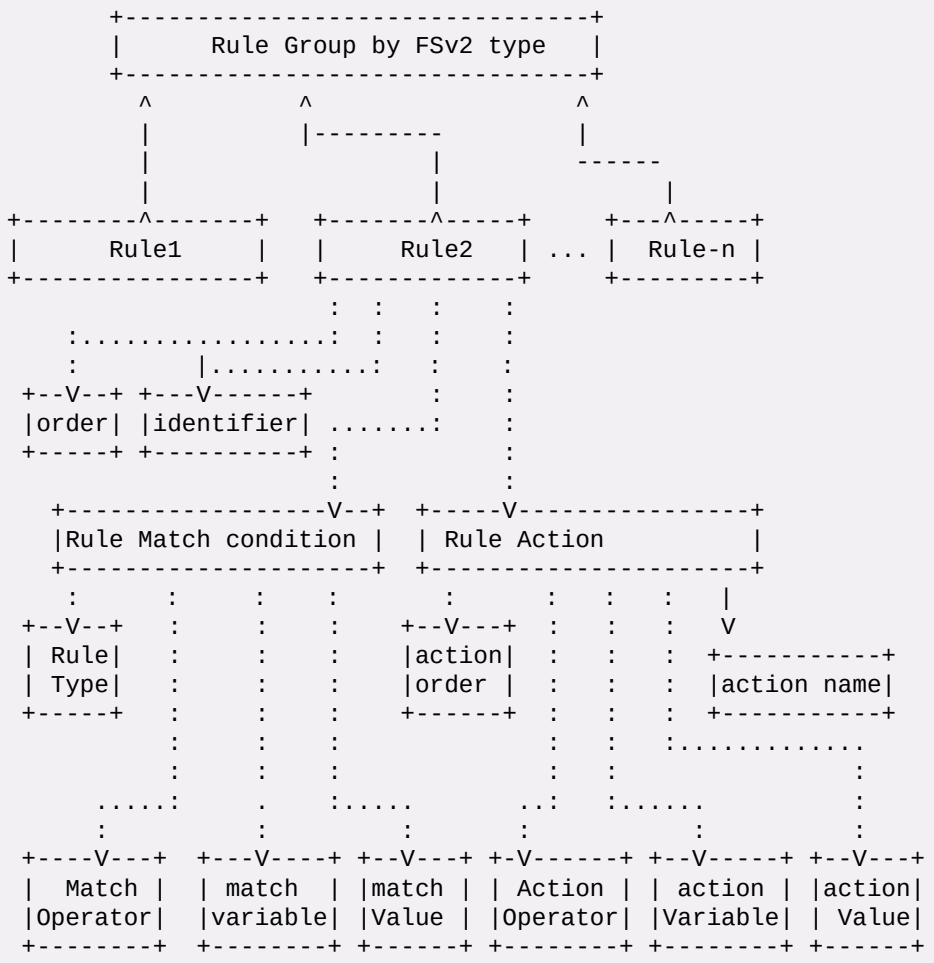

FSv2 Rules (see Rule 2) have

1. User Order – User assigned number
2. Identifier – logging identifier
3. Rule type
   - match operator
   - match variables
   - match value

Possible to have multiple filter conditions before an action

Order

Rule – 0 = permit all traffic
no actions

Rule 1-N – Filter traffic to take an action on.

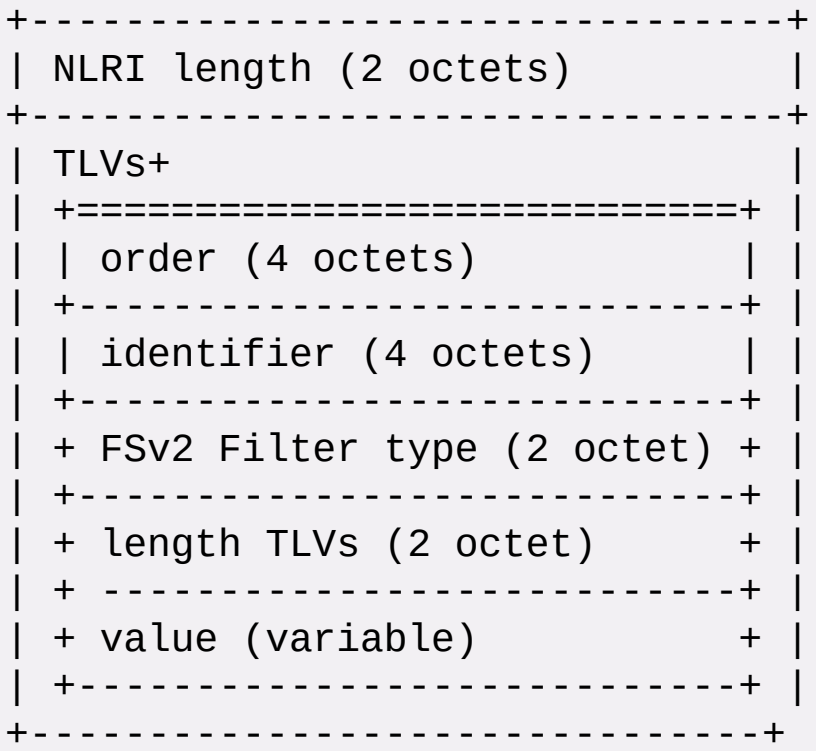

Figure 3-1 - NLRI format for FSv2

FSv2 Filter types
0 - reserved,
1 - IP Basic Filter Rules
2 - Extended IP Filter rules
3- MPLS Traffic Rules
4- L2 Traffic Rules
5- SFC Traffic rules
6 – Tunneled Traffic rules

Changed after Design Team 1 + 2 discussions

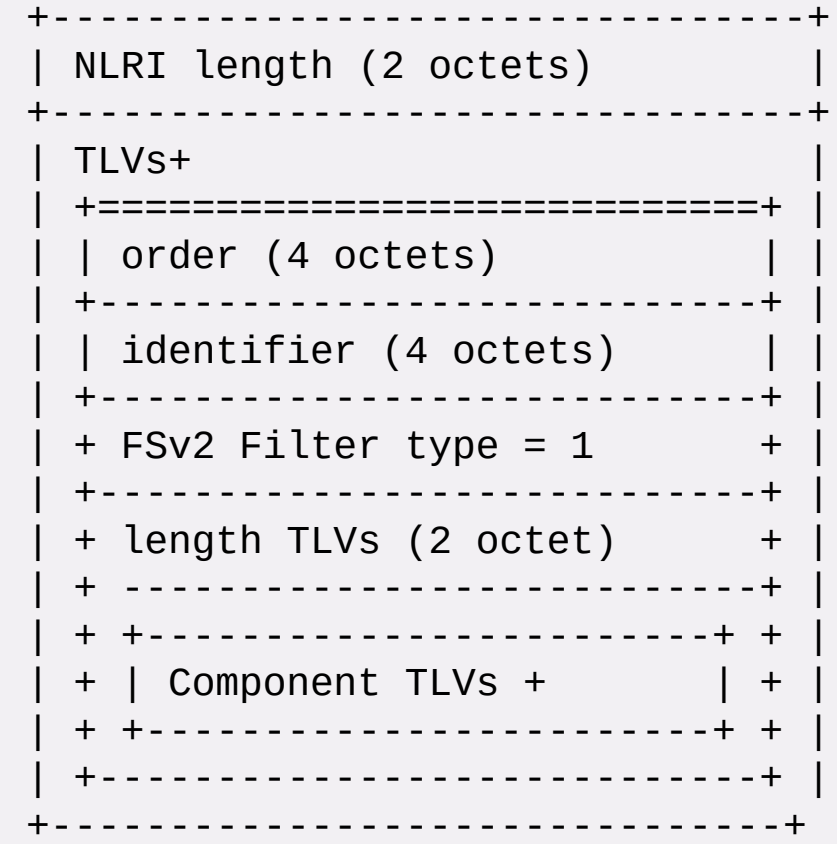

NLRI format for Basic IP Filters

Where the Component TLVs are:

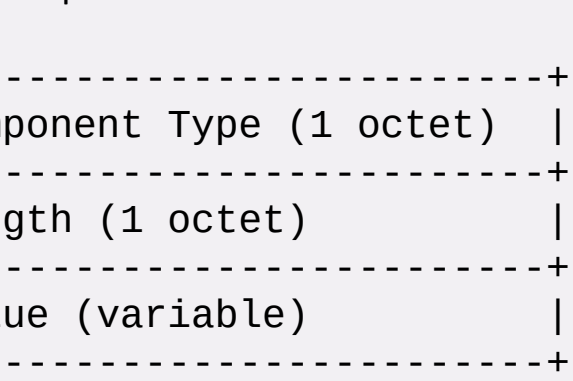

Components are
FSv1 components (IPv4, IPv6)
+ TTL field (value 14)

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

Allocation of Component IDs

15-63 Reserved for IP Extensions (standards action)
64-127 Reserved for Non-IP Filters (standards action)
128-191 Reserved (standards action)
192-249 FCFS
250-255 Reserved

IP Component numbers (in Extended IP Filters)

0/14 – TTL (option 2/option 1)
15 – SID in IPv6 Routing header
16 – NRP in Hop-by-Hop IPv6 header

Non-IP Component Numbers

64 MPLS Label Match-1
65 MPLS Label Match-2

L2 Components (AFI=25/SAFI=134)

01 L2 Ethernet
02 Source MAC
03 Destination MAC
04 DSAP
05 SSAP
06 Control in LLC
07 SNAP
08 VLAN ID
09 VLAN PCP
10 Inner VLAN ID

L2 Components (part2) (AFI=25/SAFI=134)

11 Inner VLAN PCP
12 VLAN DEI
13 Inner VALN DEI
14 Src Mac special bits
15 Dst Mac special bits

FSv2 filters + FSv1 filters

- **FSv2 and FSv1 are Ships in the night** (two NLRIs)
- **Ordering**

  Rule 0 – permit all

  Rule 1 to Rule N-1 – FSv2 with user order.

  Rule N to end - FSv1 rules at a single user order

- If same user order, then order by component number.
- If same user order + component number, then order of multiple components using rules defined in a component.

FSv2 Actions

draft-hares-idr-fsv2-ip-basic-02
draft-hares-idr-fsv2-more-ip-actions

Sue Hares

Two types of Actions

Extended Communities (EC) – FSv2-EC Actions in

- Generic transitive EC
- IPv4 Transitive Extended Communities
- Transitive IPv6-Address-Specific Actions

Community Attribute Actions

- Community Path Attribute with FSv2 Community TLV (type=2)

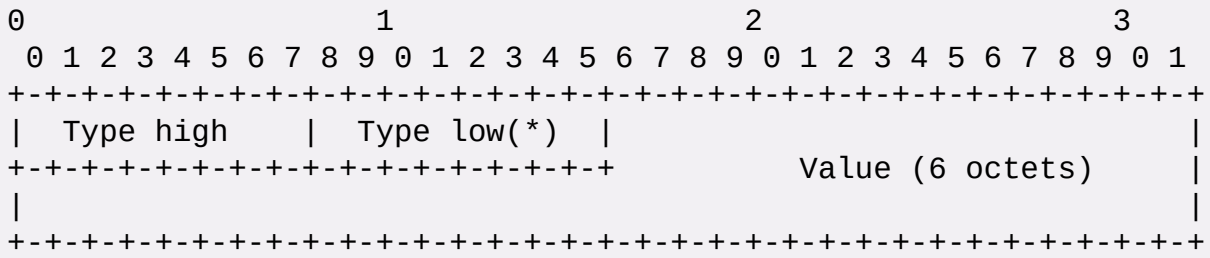

Figure 3-5 – Generic Transitive Extended Community

Table 3-3 Generic Transitive Extended Community
Part 1 - (0x80)

IPv4 Extended Communities (Type 0x80)

Value	Description	Name	Reference
0x01	**FSv2 Action Chain Ordering**	**ACO**	[This document]
0x06	FSv2 traffic-rate-byte	TRB	[RFC8955]
0x07	Flow spec traffic-action	TAIS	[RFC8955]
0x08	Flow spec rt-redirect AS-2 octet format	RDIP	[RFC8955]
0x09	Flow spec Remark DSCP	TMDS	[RFC8955]
0x0C	Flow Spec Traffic-rate-packets	TRP	[RFC8955]
0xOD	Flow Spec for SFC classifiers	SFCC	[RFC9015]

Table 3-4 Generic Transitive Extended Community
Part 2 (0x81)

IPv4 Extended Communities FSv2 action (Type 0x81)

Value	Description	Name	Reference
0x08	Flow spec rt-redirect	RDIP	[RFC8955]

Table 3-5 Generic Transitive Extended Community
Part 3 (Type 0x82)

Value	Description	Name	Reference
0x08	Flow spec rt-redirect AS-4 octet format	RDIP	[RFC8955]

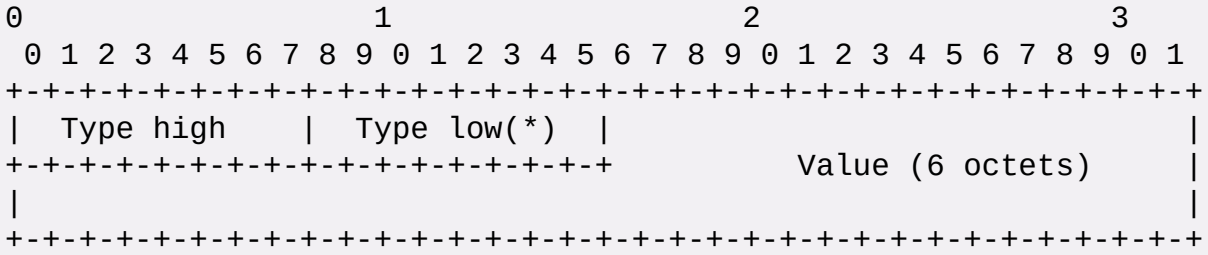

Figure 3-5 – Generic Transitive Extended Community

Table 3-7 Transitive Extended Community types (T-EC-types)

sub-type	FSv1 Description	Name	Reference
0x07	FS Interface set	Ifset	draft-ietf-idr-flowspec-interfaceset
0x08	FS Redirect/Mirror	RIPv4	draft-ietf-idr-flowspec-redirect-ip
0x09	FS Redirect to Indirection ID	RGID	draft-ietf-idr-flowspec-path-redirect

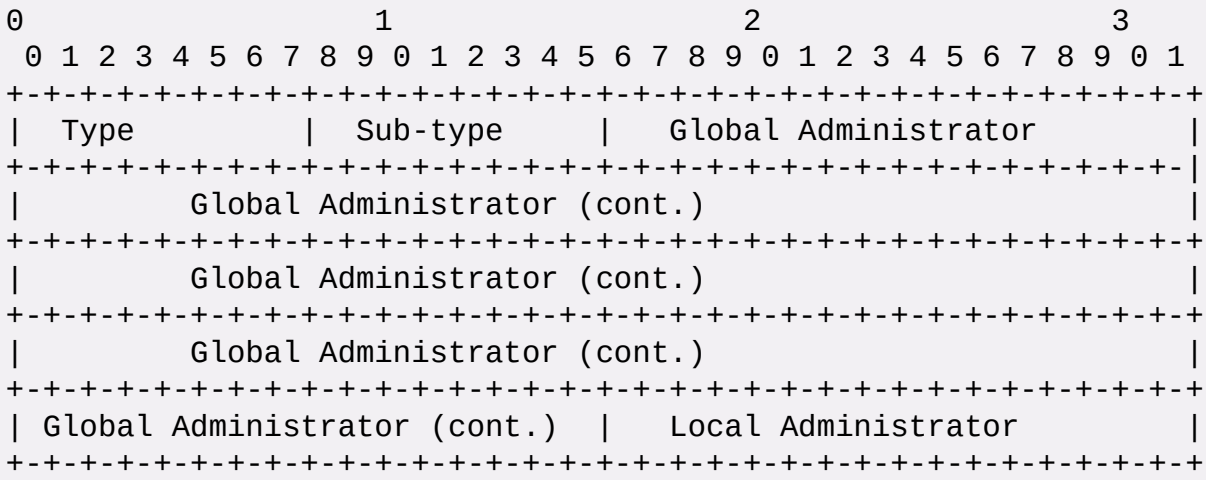

Figure 3-6 Transitive IPv6-Address-Specific-Actions

Table 3-8 Transitive IPv6-Address-Specific Actions

Value	Description	Name	
0x01	**Flow Spec Action Chain**	**ACO**	**This document**
0x0C	Flow Spec redirect-v6-flag	RD6F	draft-ietf-idr-flowspec-redirect-ip
0x0D	Flow Spec rt-redirect IPv6 format IPv6 format	RDv6	RFC8956

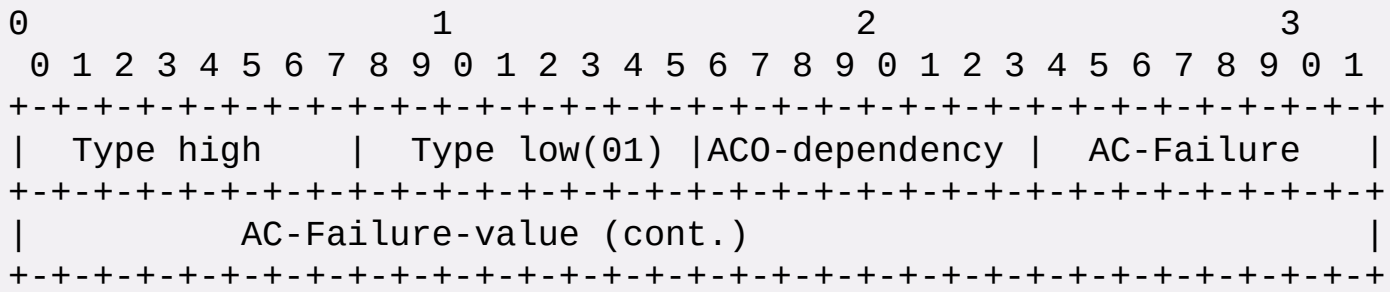

Figure 3-7

ACO Dependency - The order dependency within the Action chain.
where: 0 = default order and interactions (from this specification).
1 = Implementation specific ordering

AC Failure: Action chain action when an individual action fails
0x00 – default – stop on failure
0x01 – continue on failure (best effort on actions)
0x02 – conditional stop on failure – depending on AC-Failure-value
0x03 – rollback – do all or nothing - depending on AC-Failure-value

Yang configuration uses some of these modes on action failure.

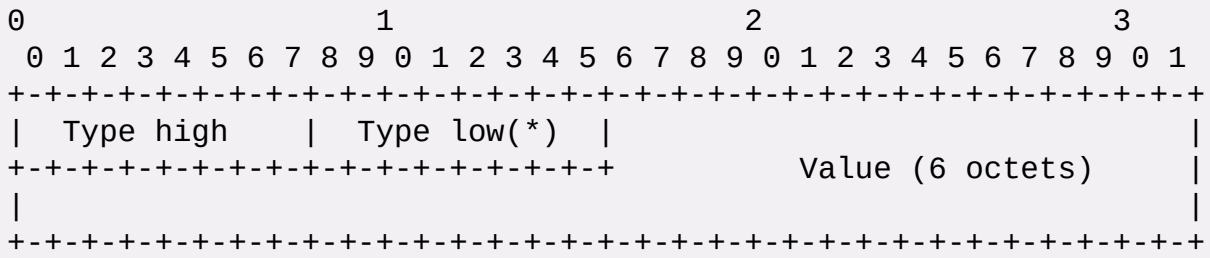

Figure 3-5 – Generic Transitive Extended Community

Table 3-3 Generic Transitive Extended Community
Part 1 - (0x80)

IPv4 Extended Communities (Type 0x80)

Value	Description	Name	Reference
0x01	**FSv2 Action Chain Ordering**	**ACO**	[This document]
0x06	FSv2 traffic-rate-byte	TRB	[RFC8955]
0x07	Flow spec traffic-action	TAIS	[RFC8955]
0x08	Flow spec rt-redirect AS-2 octet format	RDIP	[RFC8955]
0x09	Flow spec Remark DSCP	TMDS	[RFC8955]
0x0C	Flow Spec Traffic-rate-packets	TRP	[RFC8955]
0xOD	Flow Spec for SFC classifiers	SFCC	[RFC9015]
0x10	MPLS Label stack	MPLSLA	[FSv2]
TBD	VLAN Action	VLAN	[FS-l2vpn]
TBD	TPID Action	TPID	[FS-L2vpn]

Table 3-4 Generic Transitive Extended Community
Part 2 (0x81)

IPv4 Extended Communities FSv2 action (Type 0x81)

Value	Description	Name	Reference
0x08	Flow spec rt-redirect	RDIP	[RFC8955]

Table 3-5 Generic Transitive Extended Community
Part 3 (Type 0x82)

Value	Description	Name	Reference
0x08	Flow spec rt-redirect AS-4 octet format	RDIP	[RFC8955]

FSv2 in Community path header

Community Path attribute common header (figure 2-3)

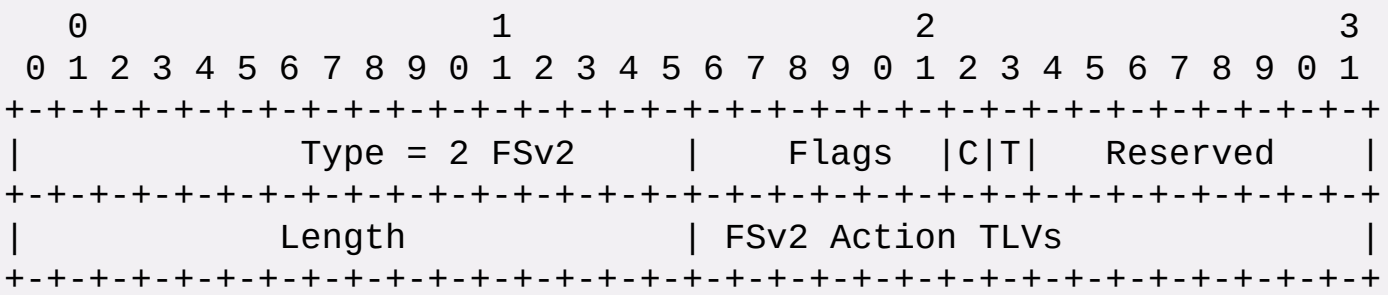

C = 1 - Transitive across Confederation boundaries
C = 0 - Non-transitive across Confederation boundaries
T = 1 – Transitive across AS boundaries
T = 0 – Non-Transitive across AS boundaries

Action TLV + SubTLV Format

Common Header for Action TLVs (figure 2-5)

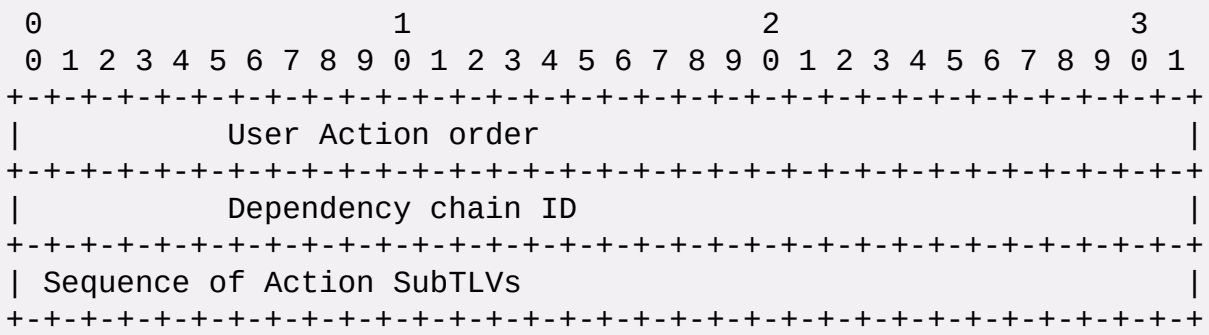

Each Action SubTLV has the format:

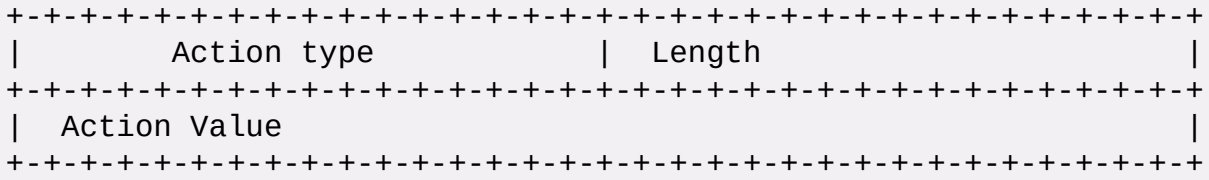

Action TLVs for Community Path Attribute

Table 5-5 All Actions Proposed for FSv2 Community Path Attribute

act-id	Name	Description	Document
TBD	MatchSet	Match and Set attribute	[IDR-rpd] (type = 03)
TBD	MatchNoA	Match and No Advertise	[IDR-rpd] (type = 04)
TBD	DetLat	Deterministic Latency action	[PD-detnet-flowmap] (type = 37)
TBD	TSNMap	Map flow to TSN stream	[PD-detnet-flowmap] (type = 38)

Action TLVs

Type 1 – Use Extended Community Actions Types + Values

- Allows Extended Community Actions with user-defined order and dependency Change

Challenge:

- What FSv2 dependency chains should there be in FSv2 actions ?
- One type of action is based on failure/success
- One type of dependency is based on a modification of packet

Nat Kao's example of modifying packet

A packet with DSCP 0 hits Rule 100.

Rule 100 has actions <Set DSCP 4, GOTO Rule 400>.

Rule 400 is matching against DSCP 4.

Will that packet be considered a match for Rule 400?

-Rule 400 will match the modified packet if we apply actions after each rule.

-Rule 400 will not match the unmodified packet, if we apply actions after all rules.

Should we modify packets as soon as the match occurs?

Design Team 3 Discussion:

4 Questions on Extended Community FSv2 Actions

Question 1: What Actions Interact?

Extended community

- Default action – is to pass all traffic.
- FSv1 default filter match action – was drop.

Action interaction in Extended Community

- Known interactions in draft-hares-idr-more-ip-actions in section 3. Are these correct?
- How do these actions interact with L2 actions or L2 detnet actions?

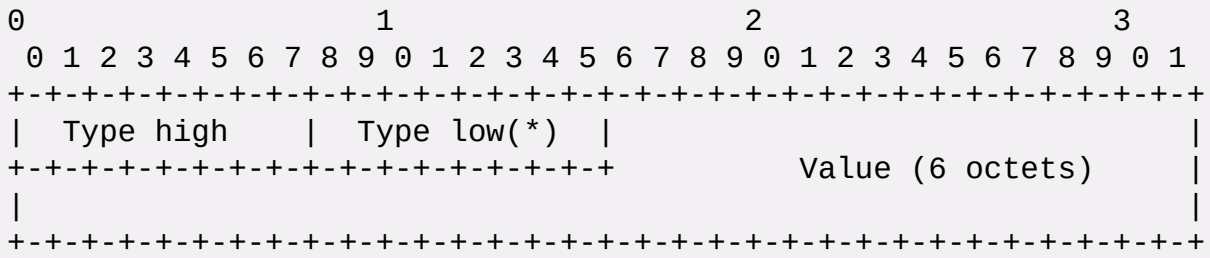

Figure 3-5 – Generic Transitive Extended Community

Table 3-3 Generic Transitive Extended Community
Part 1 - (0x80)

IPv4 Extended Communities (Type 0x80)

Value	Description	Name	Reference
0x01	**FSv2 Action Chain Ordering**	**ACO**	[This document]
0x06	FSv2 traffic-rate-byte	TRB	[RFC8955]
0x07	Flow spec traffic-action	TAIS	[RFC8955]
0x08	Flow spec rt-redirect AS-2 octet format	RDIP	[RFC8955]
0x09	Flow spec Remark DSCP	TMDS	[RFC8955]
0x0C	Flow Spec Traffic-rate-packets	TRP	[RFC8955]
0xOD	Flow Spec for SFC classifiers	SFCC	[RFC9015]
TBD	VLAN Action	VLAN	[FS-l2vpn]
TBD	TPID Action	TPID	[FS-L2vpn]

Table 3-4 Generic Transitive Extended Community
Part 2 (0x81)

IPv4 Extended Communities FSv2 action (Type 0x81)

Value	Description	Name	Reference
0x08	Flow spec rt-redirect	RDIP	[RFC8955]

Table 3-5 Generic Transitive Extended Community
Part 3 (Type 0x82)

Value	Description	Name	Reference
0x08	Flow spec rt-redirect AS-4 octet format	RDIP	[RFC8955]

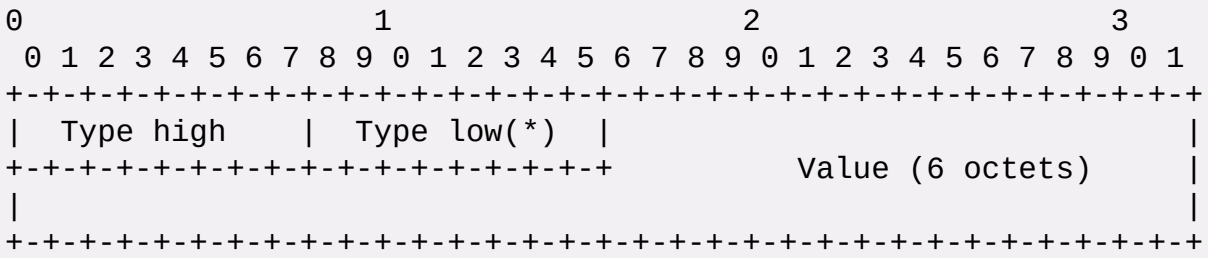

Figure 3-5 – Generic Transitive Extended Community

Table 3-7 Transitive Extended Community types (T-EC-types)

sub-type	FSv1 Description	Name	Reference
0x07	FS Interface set	Ifset	draft-ietf-idr-flowspec-interfaceset
0x08	FS Redirect/Mirror	RIPv4	draft-ietf-idr-flowspec-redirect-ip
0x09	FS Redirect to Indirection ID	RGID	draft-ietf-idr-flowspec-path-redirect

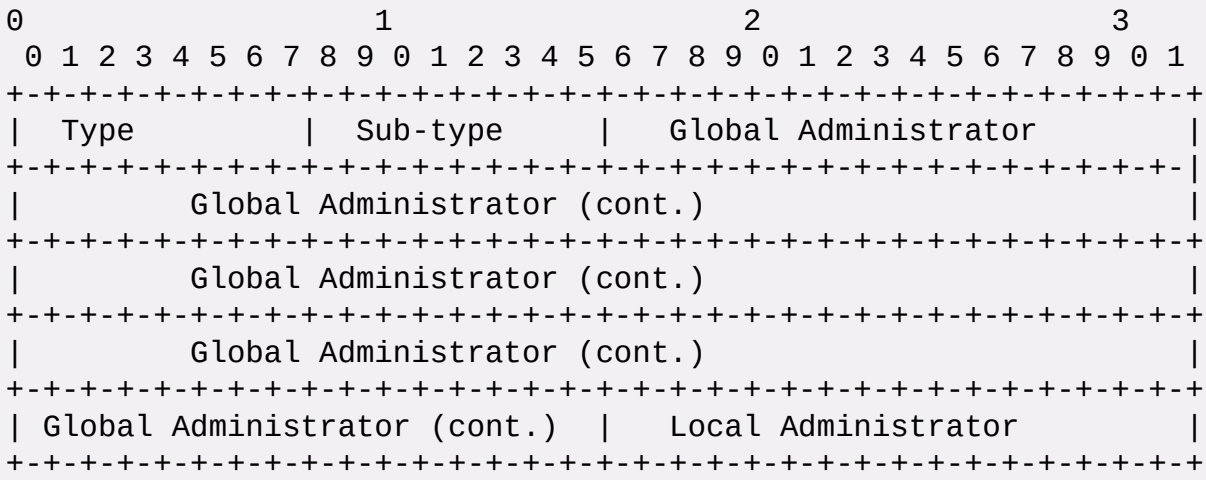

Figure 3-6 Transitive IPv6-Address-Specific-Actions

Table 3-8 Transitive IPv6-Address-Specific Actions

Value	Description	Name	
0x01	Flow Spec Action Chain	ACO	This document
0x0C	Flow Spec redirect-v6-flag	RD6F	draft-ietf-idr-flowspec-redirect-ip
0x0D	Flow Spec rt-redirect IPv6 format IPv6 format	RDv6	RFC8956

Extended Community Interactions

Value	Name	Conflicts with
0x01	ACO	none
0x06	TRB	TRP *(byte and packet traffic rate limit)*
0x07	TAIS	duplication also done in RDIP, RIPv4, RGID
0x08	RDIP	redirection done in RIPv4, RGID, copy done in TAIS
0x09	TMDS	none
0x0C	TRP	TRB
0xOD	SFCC	none

2. What Extended Community Actions do users want to order?

FSv2 has user ordering of Filters, and actions upon a match to filter.

- Should we do FSv2 default ordering of actions by action-id number for Extended Communities?
- When does the default ordering break down?
- How does the drop action get signaled?

3. What happens when Actions fail?

If an action fails, the Action Chain Ordering (ACO) FSv2-EC does:

- 0x00 - Stop on an action failure of action (default)
- 0x01 – continue on an action failure
- 0x02 – conditional stop an action failure (if value, go on)
- 0x03 – rollback – all or nothing.

For example, if a copy action fails – do you want to stop redirection?

4. What New Actions cannot be added as Extended Community

- What actions cannot be specified in the Extended Community format?
- Is it due to a lack of space for parameters?
- Or do the new actions need dependency?

Dependency

- Set the Group-ID and sub-group information?
- Does Nat’s example work without dependency?

Action TLVs for Community Path Attribute

Table 5-5 All Actions Proposed for FSv2 Community Path Attribute

act-id	Name	Description	Document
TBD	MatchSet	Match and Set attribute	[IDR-rpd] (type = 03)
TBD	MatchNoA	Match and No Advertise	[IDR-rpd] (type = 04)
TBD	DetLat	Deterministic Latency action	[PD-detnet-flowmap] (type = 37)
TBD	TSNMap	Map flow to TSN stream	[PD-detnet-flowmap] (type = 38)

Nat Kao's example of modifying packet

A packet with DSCP 0 hits Rule 100.

Rule 100 has actions <Set DSCP 4, GOTO Rule 400>.

Rule 400 is matching against DSCP 4.

Will that packet be considered a match for Rule 400?

-Rule 400 will match the modified packet if we apply actions after each rule.

-Rule 400 will not match the unmodified packet, if we apply actions after all rules.

Should we modify packets as soon as the match occurs?

Design Team 3 Discussion 2:

3 Questions on Community Attribute

user order, dependency, conflicts

5. What Actions should be implemented as User-Ordered?

- Does the DDOS finding Payload require User-ordering?
- Does better firewall rules require User-ordering?

6. What Data in actions in Community Path Attribute

Do we need:

1. Interactions with all existing actions in the Community path Attribute?
2. Interactions with all existing actions in Extended Community?
3. Ordering if same User Action-ID
4. What to do if the action fails to complete
5. Options for dependency.
6. Validity checks for parameters (parameters, duplications, and conflicts)
7. What BGP should do upon error in TLV (ignore, treat-as-withdraw)

7. How should Actions use dependency?

1. Do we have any Best Common Practices (BCP) for action dependency?
2. Any experience with dependency that failed?

Non-IP Filters – MPLS

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

Allocation of Component IDs

15-63 Reserved for IP Extensions (standards action)
64-127 Reserved for Non-IP Filters (standards action)
128-191 Reserved (standards action)
192-249 FCFS
250-255 Reserved

IP Component numbers (in Extended IP Filters)

0/14 – TTL (option 2/option 1)
15 – SID in IPv6 Routing header
16 – NRP in Hop-by-Hop IPv6 header

Non-IP Component Numbers

64 MPLS Label Match-1
65 MPLS Label Match-2

L2 Components (AFI=25/SAFI=134)

01 L2 Ethernet
02 Source MAC
03 Destination MAC
04 DSAP
05 SSAP
06 Control in LLC
07 SNAP
08 VLAN ID
09 VLAN PCP
10 Inner VLAN ID

L2 Components (part2) (AFI=25/SAFI=134)

11 Inner VLAN PCP
12 VLAN DEI
13 Inner VALN DEI
14 Src Mac special bits
15 Dst Mac special bits

MPLS label match 1 – label

type(1 octet), length(1 octet), [operator,value]+

It contains a set of {operator, value} pairs that are used for the matching filter.

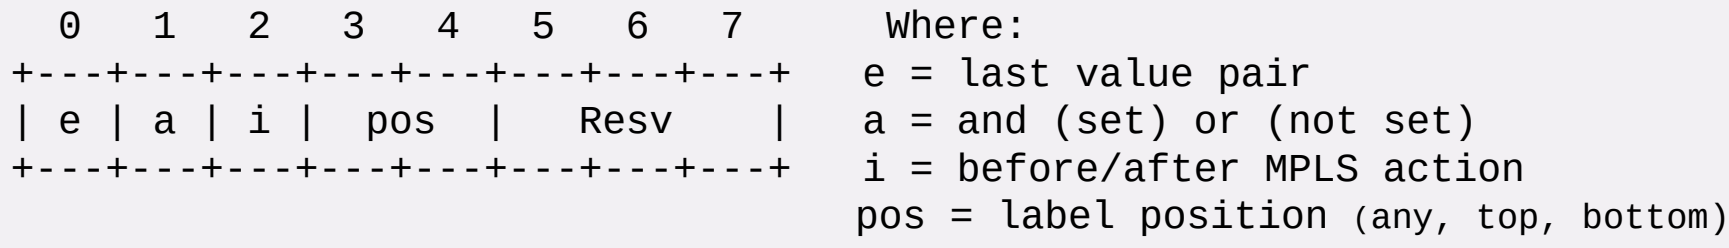

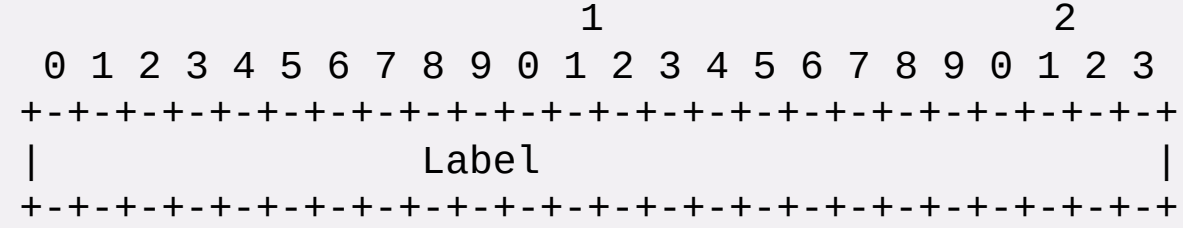

Figure 3-8

MPLS label match 2 – Experimental Bits

type(1 octet), length(1 octet), [operator,value]+

[operator, value]

- Defines a list of {operation, value} pairs used to match 3-bit exp field on the top label of packets (RFC3032)
- Values are encoded using a single byte, where the five most significant bits are zero and the three least significant bits contain the exp value.

Agenda for Design Team 4
FSv2 for Basic IP (30 minutes)

Existing Non-IP Filtering + Actions

1. MPLS + SFC – Sue Hares [5 minutes]

draft-ietf-idr-flowspec-v2, RFC9015

2. FSv1 L2VPN [15 minutes] – Donald Eastlake

draft-ietf-idr-flowspec-l2vpn

3 FSv1 Tunnels (nv03) – Donald Eastlake

draft-ietf-idr-flowspec-nv03

Non-IP Filters and Actions: L2VPN

Donald Eastlake

Non-IP Filters and Actions: Tunnels

Donald Eastlake

Thanks for all your work!