

BGP Flow Specification for Source Address Validation

[draft-geng-idr-flowspec-sav-02](#)

Nan Geng (Huawei), Dan Li (Tsinghua University)

March. 2024

BGP FlowSpec for SAV

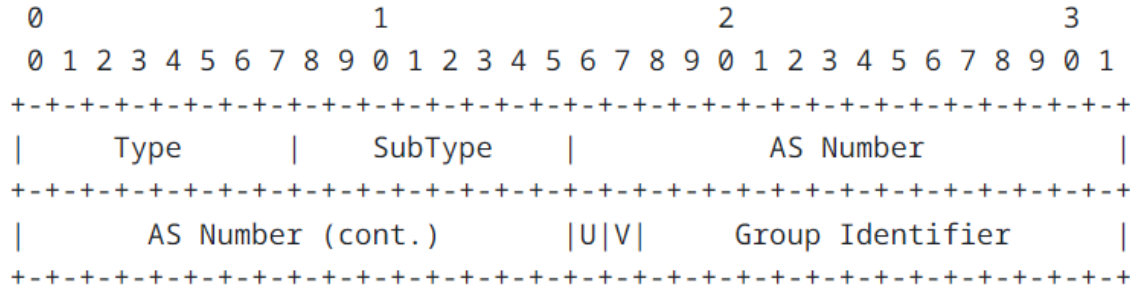
□ Background:

- ◆ **SAV rule:** <source prefix, valid-interface set> or <source prefix, invalid-interface set>.
- ◆ To facilitate SAV management and improve SAV accuracy, additional SAV rule dissemination is necessary [I-D.wu-savnet-inter-domain-architecture].
- ◆ **BGP FlowSpec** is a convenient tool for traffic filtering/controlling ([RFC8955], [RFC8956]). It supports matching source prefix but **does not support matching interfaces**.

□ This document:

- ◆ Extend BGP FlowSpec to support matching interfaces

Comments from Interim 2/26



Previous design: SAV Interface-set Extended Community (I-D.ietf-idr-flowspec-interfaceset)

- ❑ Comment 1 from Susan Hares: Interface is filter or action? I thought you are trying to do an action. if it is a filter, why put it in an extended community instead of an NLRI.
- ❑ Comment 2 from Jeffrey Haas: It is to construct the equivalence classes of interfaces that a specific source address filter will be applied to. The space you're defining here is too small even though there's a large number of bits available here. The number of equivalence classes will probably be much larger in very large networks.
- ❑ Responses:
 - ◆ Updated the draft and put interface filter in NLRI. Group ID field can be 1 to 8 Bytes.
 - ◆ Will talk with you and other people for further comments.

New Component: Incoming-Interface-Set

❑ **Encoding:** <type (1 octet), [numeric_op, value]+ >

❑ The numeric operator (numeric_op) is encoded as (see RFC8955 sec. 4.2.1.1):




Figure1

❑ **The value field is encoded as:**


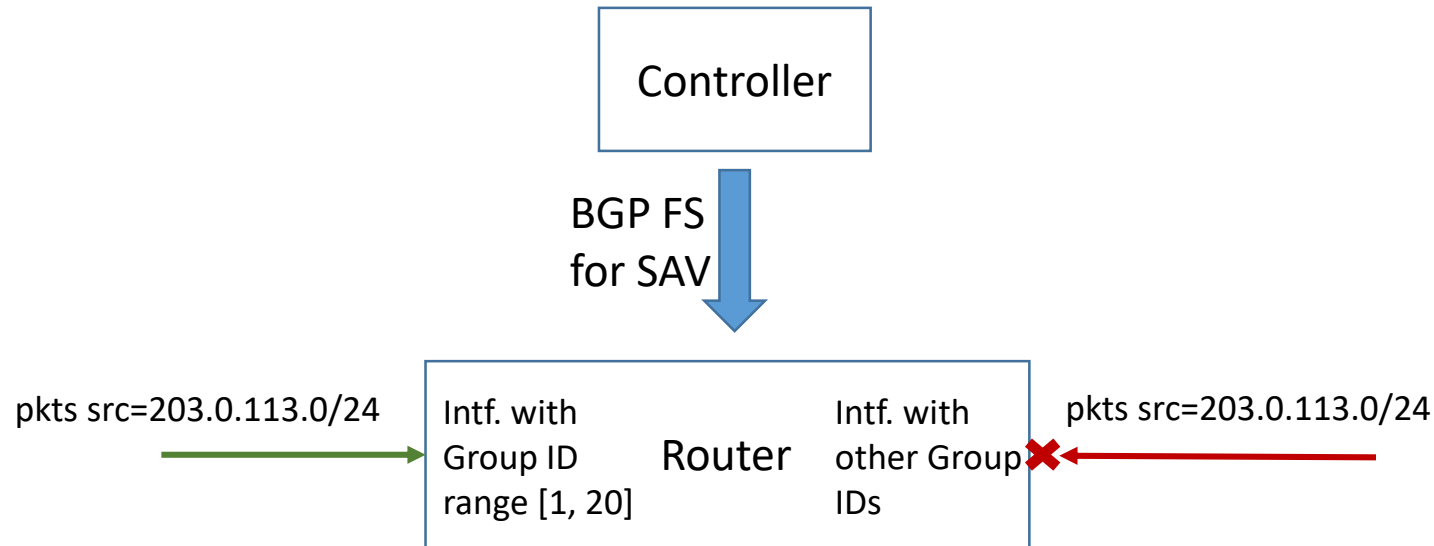


Figure2

- ❑ **Group Identifier** indicates a specific set of interfaces that are configured by network administrator.
- ❑ **Zero Group Identifier** (i.e., Group Identifier equaling 0) **any other interfaces** on the target router except the interfaces indicated by non-zero Group Identifiers in the same NLRI.
- ❑ **Group Identifier can be 1, 2, 4, or 8 octets**, which depends on the len in numeric_op.
- ❑ **Flag V (1 bit):** If set, the identified interface set is valid for source prefix. If unset, invalid.
- ❑ **Flag R (1 bit):** This bit is reserved for future use and is set to zero.
- ❑ **The bits lt, gt, and eq can be combined to match a range of Group Identifiers** (e.g., greater than Group ID1 and less than Group ID2), whose corresponding flags (i.e., V and R) **MUST** be the same.

Example

- Example: A Flow Specification NLR I encoding for "incoming interfaces {Group ID range [1, 20]} are valid for the packets from 203.0.113.0/24, and other local interfaces are invalid for the packets".



The SAV rule in BGP FS and SAV rules generated by SAV mechanisms will be put into SAV table

Length	Source	Group identifier (including flags)
0c	02 18 cb 00 71	TBD1 03 81 45 94 81 00

Value		
0x0c	length	12 octets (if len<240, 1 octet)
0x02	type	Type 2 - Source Prefix
0x18	length	24 bit
0xcb	prefix	203
0x00	prefix	0
0x71	prefix	113
TBD1	type	Type TBD1 - Incoming-interface-set
0x03	numeric_op	value size=1, >=
0x81	value	V=1, R=0, ID=1
0x45	numeric_op	"AND", value size=1, <=
0x94	value	V=1, R=0, ID=20
0x81	numeric_op	end-of-list, value size=1, ==
0x00	value	V=0, R=0, ID=0

Conclusion

- ❑ Define a Incoming-Interface-Set component (currently based on FSv1). The extended BGP FS can support matching incoming interfaces.
- ❑ SAV rule dissemination is a use case. Can also be used for other incoming interface-matching cases.
 - ◆ How to indicate the usage: R bit ?
- ❑ Comments are welcome. (On two optional designs)

Thanks!