Workgroup: IDR
Internet-Draft: draft-geng-idr-flowspec-sav-01
Published: 17 February 2024
Intended Status: Standards Track
Expires: 20 August 2024
Authors: N. Geng D. Li
Huawei Tsinghua University

BGP Flow Specification for Source Address Validation

Abstract

BGP FlowSpec reuses BGP route to distribute infrastructure and propogates traffic flow information with filtering actions. This document specifies a new BGP extended community named Source Address Validation (SAV) Interface-set to disseminate SAV rules through BGP FlowSpec.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 20 August 2024.

Copyright Notice

Table of Contents

1. Introduction

Source Address Validation (SAV) is an efficient method for preventing source address spoofing-based attacks. SAV rules indicate the valid/invalid incoming interfaces of a specific source IP address or source IP prefix. The rules can be deployed on edge routers, border routers, or aggregation routers for checking the validity of intra-domain and inter-domain packets. For invalid packets, filtering actions can be taken such as block, rate-limit, and redirect. There are many mechanisms that can generate SAV rules on routers ([RFC2827], [RFC3704], [RFC5210], [RFC8704], and [manrs-antispoofing]). However, the challenges of accurate validation and operation exist in asymetric routing scenarios or dynamic networks [I-D.ietf-savnet-intra-domain-problem-statement] [I-D.ietf-savnet-inter-domain-problem-statement]. To facilitate SAV management, additional SAV rule dissemination is needed [I-D.li-savnet-intra-domain-architecture] [I-D.wu-savnet-inter-domain-architecture].

BGP FlowSpec is a convenient and flexible tool for traffic filtering/controling ([RFC8955], [RFC8956]). It propogates traffic flow information for different traffic control purposes through the BGP protocol extension. Existing BGP FlowSpec design has supported source prefix matching and various traffic filtering actions but does not support binding valid/invalid incoming interfaces to source prefixes. With a minor extension, BGP FlowSpec can be used for SAV rule dissemination.

This document specifies a new BGP extended community named SAV Interface-set extended community. SAV rules can be disseminated

through BGP FlowSpec by combining the new extended community with source prefix component and filtering actions of existing BGP FlowSpec. The new extension can be used to configure SAV rules on remote routers. It can also act as a supplement of existing SAV mechanisms and help improve SAV accuracy.

1.1. Terminology

SAV: Source address validation

SAV Rule: The rule that indicates the valid/invalid incoming interfaces of a specific source IP address or source IP prefix.

AS: Autonomous System

1.2. Requirements Language

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Flow Specifications for SAV

2.1. SAV Rules

SAV rules can be used for checking the validity of source addresses of incoming packets. A rule usually has a format of <source prefix, interface set, validity indicator>. source prefix is for matching specific packets. Interface set represents a set of physical interfaces from which the packets arrive. Validity indicator indicates whether the packets matching the source prefix and arrival interface are valid or invalid. So, validity indicator has a value of either valid or invalid. For example, the rule <P1, [intf1, intf2], valid> means the source prefix P1 must arrive the router at interface Intf1 or Intf2, otherwise, P1 is invalid. For invalid source prefixes, the filtering actions, such as block, rate-limit, and redirect, can be taken on the packets [I-D.huang-savnet-sav-table].

In real networks, the interface set in SAV rules usually can be grouped. For example, the interfaces can be grouped as:

* Subnet interface set that contains the interfaces connecting a target subnet.

* All customer AS interfaces set or the customer AS interfaces set of a customer AS.

*All lateral peer AS interfaces set or the lateral peer AS interfaces set of a lateral peer AS.

*All transit provider AS interfaces set or the transit provider AS interfaces set of a transit provider AS.

These interface set can be indentified by a group id for easy management.

2.2. BGP FlowSpec for SAV

SAV can be disseminated to Edge/Border/Aggragation routers through BGP FlowSpec, as shown in the figure below. The controller is used to set up BGP connection with the routers in a SAV-deployed AS or domain. Note that, SAV rules disseminated by BGP FlowSpec can take effect alone or acts as a management tool of other SAV mechanisms (e.g., [RFC8704]).

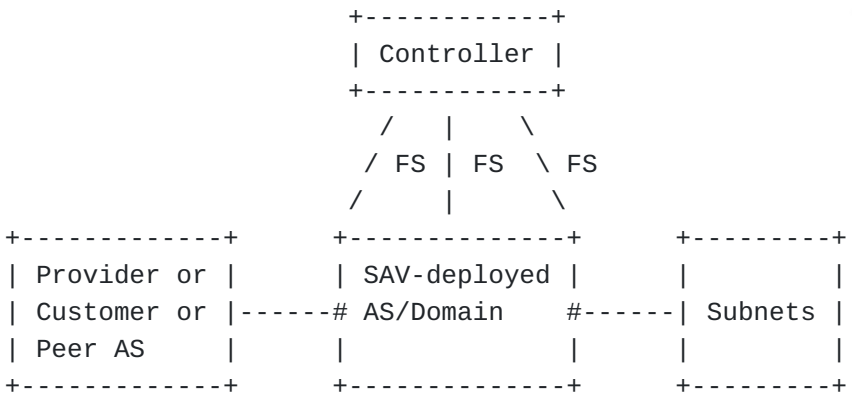

3. Extended Community for SAV

Existing BGP FlowSpec supports the component for matching source prefix and various filtering actions. This document will define a new extended community called SAV Interface-set extended community, whose design follows [I-D.ietf-idr-flowspec-interfaceset]. SAV rules can be disseminated through BGP FlowSpec by combining the new extended community with source prefix component and filtering actions of existing BGP FlowSpec ([RFC8955], [RFC8956]).

3.1. SAV Interface-set Extended Community

The newly defined SAV Interface-set extended community is encoded as follows:

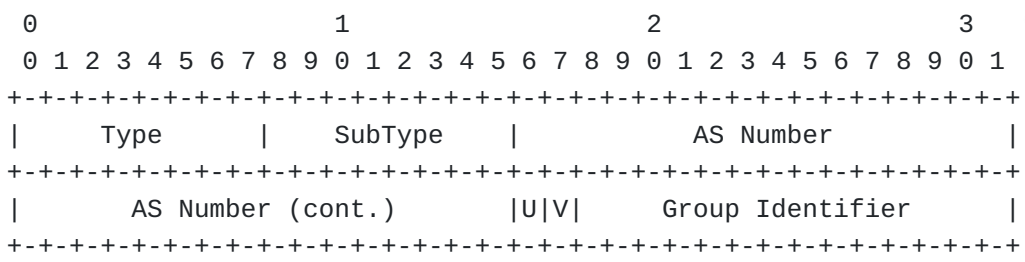

The meaning of fields:

- Type (1 octect): 0x07 or 0x47. The value of 0x07 is for FlowSpec Transitive Extended Communities, and 0x47 represents FlowSpec Non-Transitive Extended Communities. The two values have been allocated by IANA [I-D.ietf-idr-flowspec-interfaceset].

- SubType (1 octect): TBD. SubType field indicates SAV Interface-set extended community.

- AS Number (4 octects): Four-octect AS number. This field indicates the target AS where the SAV rule takes effect.

- Group Identifier (14 bits): A 14-bit number with the value ranging within 0..16383. Group identifier is a local property and identifies a set of interfaces for the source prefix carried in NLRI. The meaning of a group identifier depends on the configuration of network administrator. An interface is usually associated with one group identifiers.

- Flag V (1 bit): 1 means the identified interface set is valid for the source prefix, while 0 means the interface set is invalid for the source prefix.

- Flag U (1 bit): 1 means the rest of interfaces (not included in the interface set) on the local router are unknown for the source prefix. 0 means the rest of interfaces on the local router are invalid (when V=1) or valid (when V=0) for the source prefix.

In a BGP update, there may be more than one instances of SAV Interface-set extended community. The final interface set for the corresponding source prefix **MUST** be the union of these instances.

Multiple source prefixes can be put in multiple BGP FlowSpec NLRIs of one BGP update. In such case, these source prefixes **MUST** share the same SAV Interface-set extended communities.

3.2. Examples

Example 1: Configure soucre prefix P1 as valid at AS1's interfaces (Group Identifier=ID1) connecting a multi-homed subnet.

Encoding description: NLRI carries source prefix P1 following existing BGP FlowSpec. The SAV Interface-set community with Type=0x07 and subType=TBD carries ID1 with AS number=AS1, flag V=1, and U=1.

Example 2: Block soucre prefix P2 at AS2's interfaces (Group Identifier=ID2) connecting to transit providers.

Encoding description: NLRI carries source prefix P2 and BGP extended community carries the drop action (e.g., set traffic-rate-bytes to zero). The SAV Interface-set community with Type=0x07 and subType=TBD carries ID2 with AS number=AS2, flag V=0 and U=1.

4. IANA Considerations

This document requests a new subtype (suggested value 0x03) within the FlowSpec Transitive Extended Communities (0x07) and FlowSpec Non-Transitive Extended Communities (0x47). This sub-type shall be named "SAV Interface-set", with a reference to this document.

Value	Name	Reference
TBD	SAV Interface-set	This document

5. Security Considerations

No new security issues are introduced.

6. Acknowledgements

Many thanks to the comments from Shunwan Zhuang.

7. References

7.1. Normative References

[I-D.ietf-idr-flowspec-interfaceset] Litkowski, S., Simpson, A., Patel, K., Haas, J., and L. Yong, "Applying BGP flowspec rules on a specific interface set", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-interfaceset-05, 18 November 2019, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-interfaceset-05>.

[RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <https://www.rfc-editor.org/info/rfc8955>.

[RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <https://www.rfc-editor.org/info/rfc8956>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

7.2. Informative References

[I-D.ietf-savnet-intra-domain-problem-statement] Li, D., Wu, J., Qin, L., Huang, M., and N. Geng, "Source Address Validation in Intra-domain Networks Gap Analysis, Problem Statement, and Requirements", Work in Progress, Internet-Draft, draft-ietf-savnet-intra-domain-problem-statement-03, 13 February 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-03>.

[I-D.ietf-savnet-inter-domain-problem-statement] Wu, J., Li, D., Liu, L., Huang, M., and K. Sriram, "Source Address Validation in Inter-domain Networks Gap Analysis, Problem Statement, and Requirements", Work in Progress, Internet-Draft, draft-ietf-savnet-inter-domain-problem-statement-02, 22 August 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-problem-statement-02>.

[I-D.li-savnet-intra-domain-architecture] Li, D., Wu, J., Qin, L., Geng, N., Chen, L., Huang, M., and F. Gao, "Intra-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-li-savnet-intra-domain-architecture-06, 21 January 2024, <https://datatracker.ietf.org/doc/html/draft-li-savnet-intra-domain-architecture-06>.

[I-D.wu-savnet-inter-domain-architecture] Wu, J., Li, D., Huang, M., Chen, L., Geng, N., Liu, L., and L. Qin, "Inter-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-wu-savnet-inter-domain-architecture-06, 5 February 2024, <https://datatracker.ietf.org/doc/html/draft-wu-savnet-inter-domain-architecture-06>.

[I-D.huang-savnet-sav-table] Huang, M., Cheng, W., Li, D., Geng, N., Liu, Chen, L., and C. Lin, "General Source Address Validation Capabilities", Work in Progress, Internet-Draft, draft-huang-savnet-sav-table-03, 5 November 2023, <https://datatracker.ietf.org/doc/html/draft-huang-savnet-sav-table-03>.

[manrs-antispoofing] "MANRS Implementation Guide", January 2023, <https://www.manrs.org/netops/guide/antispoofing>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <https://www.rfc-editor.org/info/rfc3704>.

[RFC5210] Wu, J., Bi, J., Li, X., Ren, G., Xu, K., and M. Williams, "A Source Address Validation Architecture (SAVA) Testbed and Deployment Experience", RFC 5210, DOI 10.17487/RFC5210, June 2008, <https://www.rfc-editor.org/info/rfc5210>.

[RFC8704] Sriram, K., Montgomery, D., and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding", BCP 84, RFC 8704, DOI 10.17487/RFC8704, February 2020, <https://www.rfc-editor.org/info/rfc8704>.

Authors' Addresses

Nan Geng
Huawei
Beijing
China

Email: gengnan@huawei.com

Dan Li
Tsinghua University
Beijing
China

Email: tolidan@tsinghua.edu.cn