savnet
Internet-Draft
Intended status: Informational
Expires: 27 December 2025

M. Huang
Zhongguancun Laboratory
W. Cheng
China Mobile
D. Li
Tsinghua University
N. Geng
Huawei Technologies
L. Chen
Zhongguancun Laboratory

25 June 2025

General Source Address Validation Capabilities
draft-ietf-savnet-general-sav-capabilities-01

Abstract

The SAV rules of existing source address validation (SAV) mechanisms, are derived from other core data structures, e.g., FIB-based uRPF, which are not dedicatedly designed for source filtering. Therefore there are some limitations related to deployable scenarios and traffic handling policies.

To overcome these limitations, this document introduces the general SAV capabilities from data plane perspective. How to implement the capabilities and how to generate SAV rules are not in the scope of this document.

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 27 December 2025.

Table of Contents

1. Introduction

Source address validation (SAV) can detect and prevent source address spoofing on the SAV-enabled routers. When a packet arrives at an interface of the router, the source address of the packet will be validated. Invalid packets - those with unauthorized source addresses or arriving on incorrect interfaces, are typically dropped. Only validated packets will be processed or forwarded.

From the perspective of data plane validation, the SAV capabilities of existing mechanisms have two main limitations. One of them is the deployable scenario limitation. ACL rules can be configured for

filtering unauthorized source addresses at specific interfaces [RFC3704]. However, ACL is not dedicatedly designed for source prefix filtering. There exist performance and scalability issues due to long-key based searching, and typically requires expert maintenance. Strict uRPF and loose uRPF are two typical FIB-based SAV mechanisms [RFC3704] and are supported by most commercial routers/switches. FIB-based validation brings many benefits compared to ACL-based filtering but also induces some limitations. Strict uRPF is not applicable for asymmetric routing [RFC8704], which exists in various scenarios such as intra-domain multi-homing access, inter-domain interconnection, etc. Under asymmetric routing, a source prefix may have a different incoming interface from the next-hop interface of the matched entry, or the source prefix does not exist in the FIB at all. Loose mode can only block unannounced prefix, which results in massive false negatives. Overall, existing ACL-based or FIB-based SAVs can only be applied to specific scenarios and cannot be adaptive to various scenarios (e.g., symmetric vs asymmetric).

The other limitation is inflexible traffic handling policy. The current common practice is just to silently drop the spoofed packets. We don’t know who benefits from this and who is the source. Further more, the clues of attacks are ignored, which could be very helpful for dealing with DDoS attacks etc.

The root cause of the above two limitations is that there is no tool specifically designed for source address filtering. That is, the capabilities of current tools are derived from other functions, e.g., FIB or ACL.

This document describes the general SAV capabilities that the data plane of SAV-enabled devices should have. Two kinds of capabilities will be introduced: validation mode and traffic handling policy. Validation modes describe how to apply validation in different scenarios. Traffic handling policies are the policies applied on the non-validated packets. By implementing the general SAV capabilities, the above two limitations of existing mechanisms can be overcome.

To achieve accurate and scalable source address validation, dedicated SAV rules are needed instead of just using those derived from other functions, e.g., FIB or ACL.

Note that the general SAV capabilities described in this document are decoupled with vendor implementation. Conforming implementations of this specification may differ, but the SAV outcomes SHOULD be equivalent to the described SAV capabilities. And also how to generate SAV rules is not the focus of this document.

1.1. Terminology

Validation mode: The mode that describes the typical SAV application for a specific scenario. Each validation mode has its own rule syntax and validation logic.

SAV rule: The entry mapping the incoming interfaces with specific source addresses/prefixes. The SAV rule expressions and semantics might be different between validation modes.

Traffic handling policy: The policy applied to the SAV-validated 'invalid' packets.

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. Validation Modes

This section describes four validation modes (Mode1 in Section 2.1, Mode2 in Section 2.2, Mode3 in Section 2.3, Mode4 in Section 2.4). These modes take effect in different scales and need corresponding SAV rules to validate spoofing packets. By choosing modes in different scenarios appropriately, the network can be protected as much as possible while not impacting the forwarding of legitimate packets.

2.1. Mode 1: Interface-based prefix allowlist

Mode 1 is an interface-scale mode, which means it takes effect on a specific interface. The interface enabling Mode 1 is maintaining an interface-based prefix allowlist--SAV rule 1. Only the source prefixes recorded in the list will be considered valid, otherwise invalid.

Applying Mode 1 on an interface requires the complete knowledge of legitimate source prefixes connected to the interface. Mode 1 is more suitable to the closed-connected interfaces such as those connecting to a subnet, a stub AS, or a customer cone. Such a mode can efficiently prevent the connected network from spoofing source prefixes of other networks.

FIB-based strict uRPF belongs to this mode. However, to overcome the limitation of asymmetric routing, additional native source prefix-based SAV rule expression is suggested. This is essential for new SAV mechanisms or architectures such as EFP-uRPF [RFC8704], BAR-SAV [I-D.ietf-sidrops-bar-sav], Intra-domain/Inter-domain SAVNET architecture [I-D.ietf-savnet-intra-domain-architecture] [I-D.ietf-savnet-inter-domain-architecture], etc.

The scope of legitimate source prefixes for Mode 1 should ideally be as narrow and precise as possible. However, in practice due to scenario limitations, a broader scope may still be acceptable for Mode 1, as long as no legitimate prefix is omitted in the list. FIB-based loose uRPF is an extreme example of this.

Mode 1 is the most efficient one if applicable. However, in some cases, it may be difficult for an interface getting all the legitimate source prefixes. If any legitimate prefix is not included in the allowlist, packets with this source addresses arriving at the interface will be improperly blocked. For example, the interface with a default route or the interface connecting to the Internet through a provider AS can hardly promise to know all the legitimate source prefixes. We need more modes to cover those scenarios.

2.2. Mode 2: Interface-based prefix blocklist

Mode 2 is also an interface-scale mode, which means it takes effect on a specific interface. The interface enabling Mode 2 is maintaining an interface-based prefix blocklist--SAV rule 2. The source prefixes recorded in the list will be considered invalid, otherwise valid.

This mode does not require the complete knowledge of the illegitimate source prefixes on the interface. Mode 2 is suitable for proactive and reactive filtering -- Invalid source prefixes are typically preemptively added to a blocklist, enabling proactive filtering; Reactive filtering is commonly deployed by the security systems to dynamically block spoofing traffic with specific source addresses.

The prefix blocklist can be generated automatically, e.g., one of Intra-domain SAVNET architecture cases, blocking the incoming traffic with internal source prefixes on WAN interface. Or operators can configure the specific source prefixes to block from the interface, which is similar to ACL-based filtering, but more native SAV rule expression with better performance and scalability.

2.3. Mode 3: Prefix-based interface allowlist

Mode 3 is a router-scale mode, which means it can validate traffic arriving at the router from all directions. The router enabling Mode 3 will record the protected source prefixes and maintain an interface allowlist for each source prefix--SAV rule 3. If a source prefix has an interface allowlist, the packet with this source prefix is considered valid only when its incoming interface is in the interface allowlist. Otherwise, the packet is considered invalid.

Applying Mode 3 in a router requires the complete knowledge of legitimate incoming interfaces for a specific source prefix. Mode 3 focuses on validating/protecting the interested source prefixes, it is applicable to the scenario where multiple interfaces are availalbe to provide potential connection to a(or a group) specific source prefix(es) , e.g. remote AS source prefixes are connected in via the provider interfaces. Mode 3 provides a convenient and effective way to control which a group of interfaces are allowed to accept the specific source prefix, rather than to achieve similar effect by configuring Mode 2 on all other interfaces to block this source prefix.

Operators can configure the interface allowlist for a specific source prefix, to prevent DDoS attack related to this source prefix. Or the interface list for specific prefixes can be generated automatically, e.g., one capability defined by Inter-domain SAVNET architectures.

2.4. Mode 4: Prefix-based interface blocklist

Mode 4 is also a router-scale mode, which means it can validate traffic arriving at the router from all directions. The router enabling Mode 4 will maintain an interface blocklist for a specific source prefix--SAV rule 4. If a source prefix has an interface blocklist, the packet with this source prefix is considered invalid when its incoming interface is in the interface blocklist. Otherwise, the packet is considered valid.

Applying Mode 4 in a router does not requires the complete knowledge of illegitimate incoming interfaces for a specific source prefix. Mode 4 focuses on preventing specific source prefix spoofing from specific directions, it is applicable to the scenario where multiple interfaces are facing specific source prefix spoofing attack, e.g. traffic coming in a network from open connected interfaces with its internal prefix as source address. Mode 4 provides a convenient and effective way to control a group of interfaces not to accept the specific source prefix, rather than to achieve similar effect by configuring Mode 2 on each interface to block this source prefix, or Mode 3 for the specific source prefix but with a very long interface allowlist.

Operators can configure the interface blocklist for a specific source prefix, to prevent DDoS attack related to this source prefix. Or the interface list for specific prefixes can be generated automatically, e.g., one capability defined by Intra-domain SAVNET architectures.

3. Validation Procedure

Mode 1 and Mode 2 are working on interface-level, they must not be enabled on same interface at same time. If they are enabled on same interface, Mode 2 should be ignored, or be merged into Mode 1 by removing the prefix listed in Mode 2 from the allowlist of Mode 1. Mode 3 and Mode 4 are working on router-level, they are also mutual exclusive with each othter, that is, they must not be enabled for a specific source prefix at same time. If so, Mode 4 should be ignored, or be merged into Mode 3 by removing the interface listed in Mode 4 from the allowlist of Mode 3. Further more, Mode 1 are most preferred mode, which means while an interface has enabled Mode 1, the traffic for this interface don’t need go through all other modes (Mode 2, Mode 3 or Mode 4) , no matter whether they are configured. While the validation result on interface-level for Mode 2 is valid, the traffic still need go through Mode 3 or Mode 4 if applicable. Figure 1 shows a comparison of different validation modes for dealing with source address validation.

Mode	Scale	SAV rule	validation result
1	interface	1: interface-based source prefix allowlist	invalid if not matched
2	interface	2: interface-based source prefix blocklist	invalid if matched
3	router	3: prefix-based interface allowlist	invalid if not matched
4	router	4: prefix-based interface blocklist	invalid if matched

Figure 1: A comparison of different validation modes

The general validation procedure is listed as below. The final validity state, either "valid" or "invalid", will be returned after the procedure.

1) A packet arrives at the router, the source address and the incoming interface of the packet will be copied as the input for following validation process, and the initial validity state is set as ’valid’.

2) If Mode 1 is enabled on the incoming interface, the packet will be only validated based on SAV rule 1 (interface-base prefix allowlist), procedure returns with corresponding validity state. Otherwise--Mode 1 is not enabled on the incoming interface, perform following validation process.

3) If Mode 2 is enabled on the incoming interface, the packet will be validated based on SAV rule 2 (interface-base prefix blocklist). If validation result is invalid, precedure returns. If the validation result is valid or Mode 2 is not enabled, go through router-level validation procedures as below.

4) Similarly, if applicable, Mode 3 and Mode 4 validation procedure will be gone through based on SAV rule 3 (prefix-based interface allowlist) and SAV rule 4 (prefix-based interface blocklist) respectively, in which the precedure will return in case the validation result is invalid. For a specific source prefix, there should be only one router-level mode enabled.

4. Traffic Handling Policies

After doing validation, the router gets the validity state for the incoming packet. For the packet with invalid state, traffic handling policies should be executed for the packet. Simply silently dropping may not well satisfy the requirements of operators in different scenarios. This section suggests to provide flexible traffic handling policies for validated packets just like FlowSpec [RFC8955] [RFC8956].

The followings are the traffic control policies that can be taken. One and only one of the policies will be chosen for an "invalid" validation result.

* "Permit": Forward packets normally though the packets are considered invalid. This policy is useful when operators only want to monitor the status of source address spoofing in the network. Normally the "Permit" policy is configured together with the traffic monitor policies, e.g. sample.

* "Discard": Drop packets directly, which is the common choose of existing SAV mechanisms.

* "Rate limit": Enforce an upper bound of traffic rate (e.g., bps or pps) for mitigation of source address spoofing attacks. This policy is helpful while operators want do tentative filtering.

* "Traffic redirect": Redirect the packets to the specified server (e.g., scrubbing center) for attack elimination.

There are also traffic monitor policies, which are optional and can be taken together with any other policies (traffic control policies and traffix monitor policies). Some examples of the traffic monitor policies are:

* "Count": Count the number of ’invalid’ packets for each validation mode.

* "Sample": Capture the packets with a configurable sampling rate and reports them to remote servers (e.g., security analysis center) for further threat awareness and analysis.

The recommended default traffic handling policy combination is: "discard" for traffic control policy plus "count" for traffic monitor policy. The default combination could be modified per system level, per interface level, or configured based on rule level under different validation modes.

5. Relationship with Traditional SAV Mechanisms

The FIB-based SAV mechanisms (strict uPRF and loose uRPF, both belongs SAV Mode 1 -- interface based prefix allowlist) should be upgraded to the new capabilities defined in this document. By doing this, the asysmetic routing scenario limitation to strict uRPF can be overcome and new traffic handling policies can be supported, and meanwhile, the router system might not be suffering significant performance impact by doing validation based on the new SAV mechanism only, rather than on both of them.

Specially, in the network operation scenario for SAV on an open-connected interface, operator may want combine the loose uRPF and SAV Mode 2 -- loose uRPF allows only announced prefixes as source coming, and additionally SAV mode 2 blocks specific source prefixes (e.g. inner prefixes). From data plane point view, there are 2 options to address it:

1) Unified Mode 1. Maintain a prefix allowlist for the interface by deducting the source prefixes in the Mode 2 from the prefix allowlist (prefixes in FIB) in the loose uRPF.

2) Separate validation. Go through traditional loose uRPF validation first, and then go throught the Mode 2 validation.

These two options differ in aspects such as memory space organization and table lookup procedures. Option 1 is preferred if memory space in data plane is allowed.

6. Security Considerations

This document focuses on the general SAV capabilities, the generation of SAV rules is not included. There may be some security considerations for SAV generation, it is not in the scope of this document.

The "Sample" policy requires a mechanism for sampling control, sampling data encapsulation and transportation etc. The security considerations about this should be described together with the dedicated mechanism document.

7. IANA Considerations

This document includes no request to IANA.

8. Acknowledgements

The authors appreciate the valuable comments from all members of the IETF SAVNET Working Group. We extend a special thanks to Aijun Wang for his guidance and suggestion. We extend a special thanks to Mingxing Liu and Changwang Lin for their feedback and contributions from vendor implemention point of view.

9. Contributors

- Mingxing Liu

Huawei Technologies

China

Email: liumingxing7@huawei.com

- Changwang Lin

New H3C Technologies

China

email: linchangwang.04414@h3c.com

10. References

10.1. Normative References

[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>.

10.2. Informative References

[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>.

[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>.

[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>.

[I-D.ietf-sidrops-bar-sav] Sriram, K., Lubashev, I., and D. Montgomery, "Source Address Validation Using BGP UPDATEs, ASPA, and ROA (BAR-SAV)", Work in Progress, Internet-Draft, draft-ietf-sidrops-bar-sav-06, 15 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-bar-sav-06>.

[I-D.ietf-savnet-intra-domain-architecture] Li, D., Wu, J., Qin, L., Geng, N., and L. Chen, "Intra-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-ietf-savnet-intra-domain-architecture-02, 13 April 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-architecture-02>.

[I-D.ietf-savnet-inter-domain-architecture] Li, D., Chen, L., Geng, N., Liu, L., and L. Qin, "Inter-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-ietf-savnet-inter-domain-architecture-01, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-architecture-01>.

Authors' Addresses

Mingqing Huang
Zhongguancun Laboratory
Beijing
China
Email: huangmq@zgclab.edu.cn

Weiqiang Cheng
China Mobile
Beijing
China
Email: chengweiqiang@chinamobile.com

Dan Li
Tsinghua University
Beijing
China
Email: tolidan@tsinghua.edu.cn

Nan Geng
Huawei Technologies
Beijing
China
Email: gengnan@huawei.com

Li Chen
Zhongguancun Laboratory
Beijing
China
Email: lichen@zgclab.edu.cn