Workgroup: savnet
Internet-Draft:
draft-huang-savnet-sav-table-04
Published: 22 February 2024
Intended Status: Standards Track
Expires: 25 August 2024
Authors: M. Huang
Huawei Technologies
W. Cheng
China Mobile
D. Li
Tsinghua University
N. Geng
Huawei Technologies
M. Liu
Huawei Technologies
L. Chen
Zhongguancun Laboratory
C. Lin
New H3C Technologies

General Source Address Validation Capabilities

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 25 August 2024.

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. The packets with unwanted source addresses or arriving at unwanted interfaces, will be considered invalid and usually be conducted elimination actions on. Only validated packets will continue to be handled 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 unwanted 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 usually expert maintenance efforts are required. 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 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 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, a dedicated SAV table for SAV rules is needed instead of using those derived from other functions, e.g., FIB or ACL. Note that the general SAV capabilities described in this document is decoupled with real implementation. Conforming implementations of this specification may differ, but the SAV outcomes **SHOULD** be equivalent to the described SAV capabilities. How to generate SAV rules is not the focus of this document.

1.1. Terminology

SAV rule: The entry specifying the valid incoming interfaces of specific source addresses or source prefixes.

Validation mode: The mode that describes the typical applications of SAV in a specific kind of scenarios. Different modes take effect in different scales and treat the default prefix differently.

Traffic handling policy: The policy taken on the packets validated by SAV.

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 validation modes. These modes take effect in different scales and treat the default prefix differently. By choosing modes in different scenarios appropriately, the network can be protected as much as possible while not impacting the forwarding of legitimate packets.

Validation modes also describe the goal of SAV rule generation. The modes can be set before generating the rules. By specifying validation modes, operators can take appropriate SAV mechanisms matching the modes, and engineers can design new SAV mechanisms to achieve the goal in challenging scenarios.

2.1. Mode 1: Interface-based prefix allowlist

Mode 1 is an interface-scale mode, i.e., it takes effect on a specific interface. The interface enabling Mode 1 is maintaining an interface-based prefix allowlist. 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 prefixes connected to the interface. Mode 1 is 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.

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

In some cases, it may be difficult for an interface getting all the legitimate source prefixes. If not all legitimate prefixes are included in the allowlist, packets with legitimate source addresses arriving at the interface may 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.

2.2. Mode 2: Interface-based prefix blocklist

Mode 2 is also an interface-scale mode, i.e., it takes effect on a configured interface. An interface cannot enable Mode 1 and Mode 2 at the same time. The interface enabling Mode 2 is maintaining an interface-based prefix blocklist. The source prefixes recorded in the list will be considered invalid, otherwise valid.

This mode does not require the complete blocklist. If the packets with the specific source addresses need to be discarded, Mode 2 can be taken. Mode 2 is suitable for proactive filtering and reactive filtering. Usually the source prefixes that are sure to be invalid will be put into the blocklist, which is proactive filtering. Reactive filtering rules are usually installed in DDoS elimination for dropping packets 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. Or operators can configure the specific source prefixes to block from the interface. This is similar to ACL-based filtering, but more native SAV rule expression with better performance and scalability is needed.

2.3. Mode 3: Prefix-based interface list

Mode 3 is a router-scale mode, i.e., 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 list for each source prefix. The interface list of each source prefix may be an allowlist or a blocklist.

If a source prefix has an interface allowlist, the packet whose source address matches the source prefix is considered valid, only when its incoming interface is in the allowlist. Otherwise, the packet is considered invalid.

If a source prefix has an interface blocklist, the packet whose source address matches the source prefix is considered invalid, only when its incoming interface is in the blocklist. Otherwise, the packet is considered valid.

If its source address does not match any recorded source prefix, the packet is valid by default.

Mode 3 focuses on validating/protecting the interested source prefixes. Operators can configure the interface list 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 and Inter-domain SAVNET architectures.

2.4. Validation Procedure

Mode 1 and Mode 2 are working on interface-level, while Mode 3 is for the router-level. Thus, there can be multiple modes configured on the same router. Mode 1 are most preferred if applicable (with best pretection effect) and mutual exclusive with the other two modes, which means while an interface enabled Mode 1, the traffic for this interface don't need go through Mode 2 or Mode 3 -- while an interface enabled Mode 2, the traffic still need go through Mode 3. Figure 1 shows a comparison of different validation modes for dealing with default prefix.

Mode	Scale	Treatment of the default prefix
1	interface	invalid
2	interface	valid
3	router	check its incoming interface

Figure 1: A comparison of different validation modes

The validation procedure is shown in Figure 2. Suppose the router has learned the SAV rules by SAV mechanisms and implemented them in the data plane. When a packet arrives at the router, the router will take the source address and the incoming interface of the packet as the input and look up the SAV rules. The final validity state that is either "valid" or "invalid" will be returned after the procedure.

Firstly, the packet is validated by the enabled interface-scale mode, i.e., Mode 1 or Mode 2. If the validity state1 or validity state2 is "invalid", the final validity state is "invalid" and the packet does not have to be validated by Mode 3. If the validity state1 or validity state2 is "valid", the packet still needs to be validated by the enabled Mode 3 and the validity state3 is the final validity state.

If a mode is not enabled, then the corresponding list should not be queried. If Mode 3 is not enabled, either the validity state1 or the validity state2 is the final validity state.

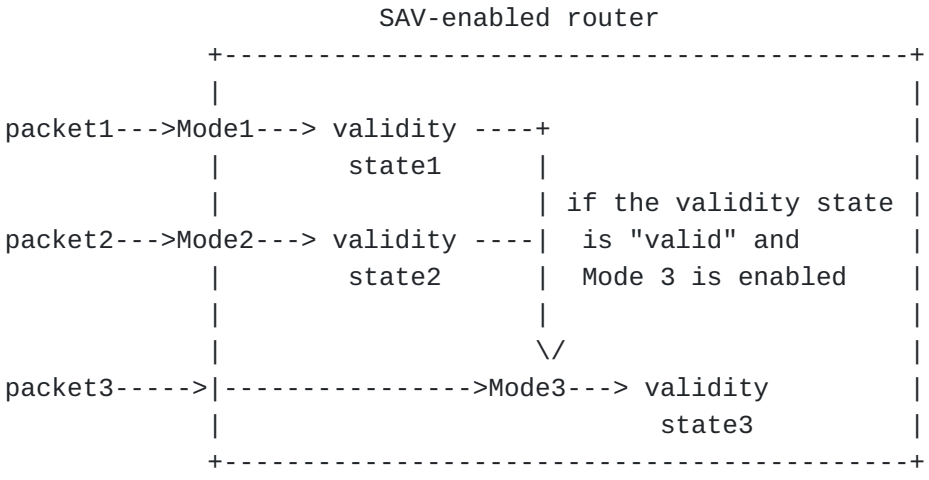

Figure 2: Validation procedure

To achieve accurate and scalable source address validation, a dedicated SAV table for SAV rules is needed rather than using those derived from other functions, e.g., FIB or ACL.

3. Traffic Handling Policies

After doing validation, the router gets the validity state of the incoming packet. For the packet with invalid state, traffic handling policies should be taken on the packet. Simply forwarding or silently dropping may not well satisfy the requirements of operators in different scenarios. This section suggests to provide flexible traffic handling policies to 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. The "Permit" policy can be taken together with the "Sample" policy.

* "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 points (e.g., scrubbing centers) in the network for attack elimination.

There are also traffic monitor policies that are optional. One of the useful traffic monitor policies is:

*"Sample": Capture the packets with a configurable sampling rate and reports them to remote servers (e.g., security analysis center). The sampled packets can be used for threat awareness and further analysis [I-D.cheng-savnet-proactive-defense-network]. "Sample" can be taken together with any one of the above policies. Note that, existing techniques like NetStream or NetFlow can be used for "Sample".

4. Security Considerations

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

The "Sample" policy pushes data to remote servers. This function can be achieved by existing techniques like NetStream or NetFlow. The "Sample" policy may induce same security considerations as these techniques, and the corresponding documents have discussed them.

5. IANA Considerations

This document includes no request to IANA.

6. References

6.1. 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>.

6.2. 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.cheng-savnet-proactive-defense-network] Cheng, W., Geng, N., Li, D., and Yue, "Network Proactive Defense based on Source Address Validation", Work in Progress, Internet-Draft, draft-cheng-savnet-proactive-defense-network-01, 18 October 2023, <https://datatracker.ietf.org/doc/html/draft-cheng-savnet-proactive-defense-network-01>.

[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-02, 12 September 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-bar-sav-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>.

Authors' Addresses

Mingqing Huang
Huawei Technologies
Beijing
China

Email: huangmingqing@huawei.com

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

Mingxing Liu
Huawei Technologies
Beijing
China

Email: liumingxing7@huawei.com

Li Chen
Zhongguancun Laboratory
Beijing
China

Email: lichen@zgclab.edu.cn

Changwang Lin
New H3C Technologies
Beijing
China

Email: linchangwang.04414@h3c.com