

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 15 July 2022

D. Li
J. Wu
Tsinghua University
M. Huang
Huawei
L. Qin
Tsinghua University
N. Geng
Huawei
11 January 2022

Source Address Validation: Use Cases and Gap Analysis
draft-li-sav-gap-analysis-01

Abstract

This document identifies the importance and use cases of source address validation (SAV) at both intra-domain level and inter-domain level (see [\[RFC5210\]](#)). Existing intra-domain and inter-domain SAV mechanisms, either Ingress ACL filtering [\[RFC2827\]](#), unicast Reverse Path Forwarding (uRPF) [\[RFC3704\]](#), or Enhanced Feasible-Path uRPF (EFP-uRPF) [\[RFC8704\]](#) has limitations in scalability or accuracy. This document provides gap analysis of the existing SAV mechanisms.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 8174](#) [\[RFC8174\]](#).

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 15 July 2022.

Internet-Draft

SAV Use Case & Gap Analysis

January 2022

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	4
3.	Use Cases	4
3.1.	Use Case 1: Intra-domain SAV	5
3.2.	Use Case 2: Inter-domain SAV	6
4.	Gap Analysis	6
4.1.	Existing Intra-domain SAV mechanisms	6
4.2.	Existing Inter-domain SAV mechanisms	7
5.	SAV Requirements	9
6.	Security Considerations	9
7.	Contributors	9
8.	Acknowledgments	9
9.	Normative References	9
	Authors' Addresses	10

[1.](#) Introduction

Source Address Validation (SAV) is important for defending against source address forgery attacks and accurately tracing back to the attackers. Considering that the Internet is extremely large and complex, it is very difficult to solve the source address spoofing problem at a single "level" or through a single SAV mechanism. On the one hand, it is unrealistic to require all networks to deploy a single SAV mechanism. On the other hand, the failure of a single SAV mechanism will completely disable SAV.

To address the issue, Source Address Validation Architecture (SAVA) was proposed [[RFC5210](#)]. According to the operating feature of the Internet, SAVA presents a hierarchical architecture which carries out source IP address validation at three checking levels, i.e., access network, intra-domain, and inter-domain. Different levels provide

different granularities of source IP address authenticity. In contrast to the single-level/point model, SAVA allows incremental deployment of SAV mechanisms while keeps effective because of its multiple-fence design. So, enhancing the source IP address validity in all the three checking levels is of high importance. Furthermore, one or more independent and loosely-coupled SAV mechanisms can coexist and cooperate under SAVA, which is friendly to different users (e.g., providers) with different policies or considerations. Obviously, the quality of SAV mechanisms for their target checking levels is key to the performance of SAV.

There are many SAV mechanisms for different checking levels. For the access network level, Source Address Validation Improvement (SAVI) was proposed to force each host to use legitimate source IP address [[RFC7039](#)]. SAVI acts as a purely network-based solution without special dependencies on hosts. It dynamically binds each legitimate IP address to a specific port/MAC address and verifies each packet's source address through the binding relationship. One of the most attractive features of SAVI is that it supports the maximally fine granularity of individual IP addresses, which previous ingress filtering mechanisms cannot provide.

At the intra-domain level, static Access Control List (ACL) is a typical solution of SAV. Operators can configure some matching rules to specify which kind of packets are acceptable (or unacceptable). The information of ACL should be updated manually so as to keep consistent with the newest filtering criteria, which inevitably limits the flexibility and accuracy of SAV. Strict unicast Reverse Path Forwarding (uRPF) [[RFC3704](#)] is another solution suitable to intra-domain. Routers deploying strict uRPF accept a data packet only when i) the local FIB contains a prefix encompassing the packet's source address and ii) the corresponding forwarding action for the prefix matches the packet's incoming interface. Otherwise, the packet will be dropped. However, in the scenarios (e.g., multihoming cases) where data packets are under asymmetric routing, strict uRPF often improperly blocks legitimate traffic.

At the inter-domain level, a combination of Enhanced Feasible-Path uRPF (EFP-uRPF) and loose uRPF is recommended in[RFC8704]. Particularly, EFP-uRPF is suggested to be applied on customer interfaces. EPF-uRPF on an AS can prevent its customers from spoofing its upstream ASes' source addresses but fails in the case of two customers spoofing each other. On lateral peer interfaces and transit provider interfaces, loose uRPF [RFC3704] is taken. The routers deploying loose uRPF accept any packets whose source addresses appear in the local FIB tables. Due to the loss of directionality, loose uRPF often improperly permits spoofed traffic.

To summarize, given that it is impossible to deploy SAVI on every access network in the Internet, the "fences" at intra- and inter-domain levels are very important for filtering source address forgery packets that are let go by access networks. However, there exist some instinctive drawbacks in the existing SAV mechanisms designed for both the intra- and inter-domain levels, which leads to inevitable improper permit or improper block problems. A more complete SAV mechanism is required for both intra- and inter-domain levels.

This document identifies the use cases of intra- and inter-domain SAVs. These cases will help analyze the instinctive drawbacks of the existing SAV mechanisms. After that, some SAV requirements will be presented.


[2.](#) Terminology

EAST-WEST traffic denotes the traffic originated and terminated within an AS. Intra-domain SAV aims to check EAST-WEST traffic and prevents hosts/routers from spoofing other source IP blocks in the same AS.

NORTH-SOUTH traffic denotes the traffic arriving from an external AS. Particularly, the traffic arriving from the customer AS is Northward traffic. The traffic received from the provider/peer AS is Southward traffic. Inter-domain SAV aims to verify the authenticity of the source address of NORTH-SOUTH traffic.

[3.](#) Use Cases

Figure 1 illustrates the use cases of SAV in both intra- and inter-domain levels. AS1-AS5 belong to the same customer cone, and AS1 is the stub AS. The topology of AS2 is presented while other ASes' inner structures are hidden for brevity.



+-----+

P1 is the source IP address prefix of Router3.

P1' is the spoofed P1 by Router2 located in the same AS as Router3.

P1" is the spoofed P1 by Routers located in another AS, i.e., AS4.

Figure 1: Illustration of the use cases of SAV in both intra- and inter-domain levels

[3.1.](#) Use Case 1: Intra-domain SAV

In some scenarios especially very large ASes, hosts/routers in the same AS may spoof each other's IP addresses. In Figure 1, Router2 spoofs P1 that originates from Router3. With Intra-domain SAV, EAST-WEST traffic can be checked, and source address spoofing attacks can be prevented. In the figure, Router1, Router3, and Router4 will drop the packets with P1' while accept those with P1, when they deploy Intra-domain SAV mechanisms. Overall, Intra-domain SAV can prevent the source address spoofing from the same AS.

[3.2.](#) Use Case 2: Inter-domain SAV

In Figure 1, AS4 spoofs AS2's IP address prefix, i.e., P1 originated from Router3. AS5 will receive the Northward traffic from AS2 and AS4 with legitimate and spoofed IP addresses, respectively. An SAV mechanism is necessary for AS5 to drop the illegal traffic. From the view point of Southward traffic, AS1 may also receive spoofed traffic from AS3 (if AS3 accepts the data packets with source prefix P1"). So, the deployment of SAV on AS1 is also important. Overall, Inter-domain SAV is necessary and can improve the confidence of the source IP address validity among ASes.

[4.](#) Gap Analysis

High accuracy is the basic requirement of any intra- or inter-domain SAV mechanism. For any SAV mechanism, improper block problems must be avoided because legitimate traffic must not be influenced. On that basis, SAV should also reduce improper permit problems as much as possible. However, existing SAV mechanisms can not well meet these requirements.

4.1. Existing Intra-domain SAV mechanisms

Operators can configure static ACLs on border routers to validate source addresses. The main drawback of ACL-based SAV is the high operational overhead. Limited application scenarios make the ACL-based method unable to do sufficient SAV on EAST-WEST traffic.

Strict uRPF can generate SAV tables automatically, but it also has limited application scenarios. Figure 2 illustrates an intra-domain scenario. In the scenario, AS1 runs strict uRPF. An access network having IP address prefix 10.0.0.0/15 is attached to two border routers (Router1 and Router2) of AS1. Due to customer's policy, it advertises 10.0.0.0/16 to Router1 and 10.1.0.0/16 to Router2. Then, Router1 and Router2 will advertise the learned IP address prefixes to other routers in AS1 through intra-domain routing protocols such as OSPF and IS-IS.

Although customer only advertises 10.0.0.0/16 to Router1, it may send packets with source IP addresses belonging to 10.1.0.0/16 to Router1 due to load balancing requirements. Suppose the destination node is Router5. Then the path to destination is Customer->Router1->Router3->Router5, while the reverse path is Router5->Router4->Router2->Customer. The round trip routing path is asymmetric, which cannot be dealt with well by strict uRPF.

Specifically speaking, strict uRPF is faced with improper block problems under asymmetric routing scenarios. When Router1/Router3 runs strict uRPF, it learns SAV rules that packets with source address prefix of 10.0.0.0/16 must enter the router on interface '#'. When the packets with source addresses of 10.1.0.0/16 arrive, they will be dropped, which results in improperly blocking legitimate traffic. Similarly, when strict uRPF is deployed on Router2, the improper block problem still exists.




Figure 2: An intra-domain scenario

4.2. Existing Inter-domain SAV mechanisms

The most popular inter-domain SAV is suggested by [RFC8704], which combines EFP-uRPF algorithm B and loose uRPF. In particular, EFP-uRPF algorithm B is for Northward traffic validation. It sacrifices the directionality of customer interfaces for reducing improper permit cases. Loose uRPF is for validating Southward traffic on lateral peer and transit provider interfaces. It sacrifices directionality of Southward traffic completely. Such a combined method sacrificing directionality will leads to improper permit problems sometimes.

Figure 3 illustrates a common inter-domain scenario where the above inter-domain SAV method will fail. In the figure, there are two customer ASes, i.e., AS1 and AS2. Both of them are attached to a

provider AS, i.e., AS4. AS4 has a lateral peer and a provider, i.e.,

AS3 and AS5. Particularly, AS1 has IP address prefix P1 and advertises it to AS4. IP address prefix P2 is allocated to AS2 and is also advertised to AS4. AS3 has IP address prefix P3 and AS5 has IP address prefix P5. P3 and P5 are also advertised to AS4 through BGP. All arrows represent BGP advertisements. Assume AS4 deploys inter-domain SAV policies, i.e., a combination of EFP-uRPF algorithm B and loose uRPF.

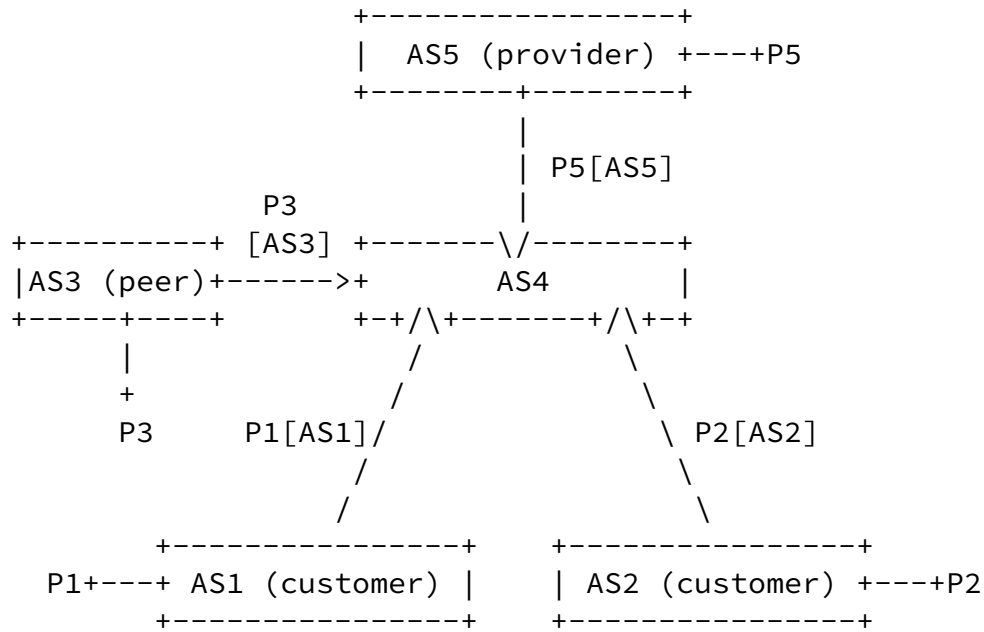


Figure 3: An inter-domain scenario

For Northward traffic, AS4 applies EFP-uRPF. Under EFP-uRPF, AS4 will generate SAV rules considering P1 and P2 are legitimate on both the two customer interfaces. When AS1 spoofs IP address prefix P2 of AS2, the malicious Northward traffic cannot be filtered by AS4. The same is true when AS2 forges P1 of AS1. That is to say, EFP-uRPF cannot prevent source address spoofing among customers even though it only focus on Northward traffic.

For Southward traffic, AS4 deploys loose uRPF for the interfaces of AS3 and AS5. It will learn that the packets with source addresses of P3 or P5 can be accepted without validating the specific arrival interface. Since loose uRPF loses directionality completely, it obviously will fail in dealing with the source address spoofing between its lateral peer and provider, i.e., AS3 and AS5.

[5.](#) SAV Requirements

High accuracy, i.e. avoiding improper block problems while trying best to reduce improper permit problems, is the basic requirement of an ideal SAV mechanism. As described above, existing SAV mechanisms cannot meet this requirement. The root cause of their limitations is that they all achieve SAV based on local forwarding information base (FIB) or routing information base (RIB), which may not match the real forwarding direction from the source. In order to guarantee the accuracy, SAV should follow the real data-plane forwarding path. To solve this problem and provide accurate SAV for arbitrary network scenarios, it is required to exchange/explore/probe the forwarding-path information among routers/ASes. In other words, network-wide protocols should be considered.

The network-wide protocols should also consider some practical issues:

- * High scalability. The protocols should not induce much overhead (e.g., bandwidth cost of path probing). Fast convergence under environment changes is also important for improving the scalability in different scales of networks.
- * High deployability. A strategy of incremental deployment needs to be considered. If some routers/ASes do not support the new protocols, improper block should be avoided.
- * High security. The protocols should include mechanisms to guarantee the integrity of protocol packets. Security risks such as Man-in-the-Middle Attack should be avoided.

[6.](#) Security Considerations

TBD

[7.](#) Contributors

TBD

[8.](#) Acknowledgments

TBD

[9.](#) Normative References

Internet-Draft

SAV Use Case & Gap Analysis

January 2022

- [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>>.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", [RFC 7039](#), DOI 10.17487/RFC7039, October 2013, <<https://www.rfc-editor.org/info/rfc7039>>.
- [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>>.
- [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

Dan Li
Tsinghua University
Beijing
China

Email: tolidan@tsinghua.edu.cn

Jianping Wu
Tsinghua University
Beijing
China

Email: jianping@cernet.edu.cn

Li, et al.

Expires 15 July 2022

[Page 10]

Internet-Draft

SAV Use Case & Gap Analysis

January 2022

Mingqing Huang
Huawei
Beijing
China

Email: huangmingqing@huawei.com

Lancheng Qin
Tsinghua University
Beijing
China

Email: qlc19@mails.tsinghua.edu.cn

Nan Geng
Huawei
Beijing
China

Email: gengnan@huawei.com

