Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 19 August 2026

D. Li
Tsinghua University
L. Qin
L. Liu
Zhongguancun Laboratory
M. Huang
Huawei
K. Sriram
USA NIST

15 February 2026

Gap Analysis, Problem Statement, and Requirements for Inter-Domain SAV
draft-ietf-savnet-inter-domain-problem-statement-14

Abstract

This document provides a gap analysis of existing inter-domain source address validation mechanisms, describes the problem space, and defines the requirements for technical improvements.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Source Address Validation (SAV) is a fundamental mechanism for detecting and mitigating source address spoofing attacks [RFC2827] [RFC5210] [RFC3704] [RFC8704]. This document provides a gap analysis of existing inter-domain SAV mechanisms, describes the problem space, and defines the requirements for technical improvements. The corresponding work related to intra-domain SAV is documented in [I-D.ietf-savnet-intra-domain-problem-statement].

In this document, inter-domain SAV refers to SAV on AS-to-AS interfaces that carry external BGP (eBGP) sessions. The eBGP sessions include Customer-to-Provider (C2P), Provider-to-Customer (P2C), lateral peering (p2p), and Route Server to RS-client connections. The terms customer, provider (transit provider), and lateral peer (non-transit peer; peer (for simplicity)) used in this document are consistent with those defined in [RFC7908] [RFC9234]. Further, [RFC9234] mentions Route Server (RS) and RS-client. An RS-to-RS-client interface is akin to the customer interface. For the purposes of SAV, an RS-client-to-RS interface may be treated (1) like a provider interface for simplicity, or (2) like a union of lateral peers considering all the ASes the RS-client chose to peer with at the IXP RS.

Access Control List (ACL) and unicast Reverse Path Forwarding (uRPF) based techniques are currently utilized to some extent for inter-domain SAV. In this document, the inter-domain SAV methods from only the existing IETF RFCs (BCP 38 [RFC2827] and BCP 84 [RFC3704] [RFC8704]) are considered for the gap analysis; IETF work-in-progress documents are not considered. This document analyzes the available methods and attempts to answer: (1) what are the technical gaps (Section 4), (2) what are the outstanding problems (problem statement) (Section 5), and (3) what are the practical requirements for the solutions to these problems (Section 6).

The following summarizes the fundamental problems with existing SAV mechanisms, as analyzed in Section 4 and Section 5:

* Improper block: Existing uRPF-based mechanisms suffer from improper block (false positives) in two inter-domain scenarios: limited propagation of a prefix and hidden prefix.

* Improper permit: With some existing uRPF-based SAV mechanisms, improper permit (false negatives) can happen on any type of interface (customer, lateral peer, or provider). Specifically, if the method relaxes the directionality constraint [RFC3704] [RFC8704]} to try to achieve zero improper blocking, the possibility of improper permit increases. (Note: It is recognized that unless there is full adoption of SAV in the customer cone (CC) of the interface in consideration, improper permit is not fully preventable in scenarios where source address spoofing occurs from within the CC, i.e., a prefix at one Autonomous System (AS) in the CC is spoofed from another AS in the same CC.)

* High operational overhead (HOO): ACL-based ingress SAV filtering introduces significant operational overhead, as it needs to update ACL rules manually to adapt to prefix or routing changes in a timely manner. The HOO issue does not pertain to existing uRPF-based mechanisms.

To address these problems, this document specifies (Section 6) the following key technical requirements for any new solution:

* Improved SAV accuracy over existing mechanisms: Any new inter-domain SAV mechanism MUST avoid improper blocking and have superior directionality property (reject more spoofed traffic) than existing inter-domain SAV mechanisms.

* Reduced operational overhead: Any new inter-domain SAV mechanism MUST be able to automatically adapt to network dynamics and asymmetric routing scenarios. Any such mechanism MUST have less operational overhead than ACL-based ingress SAV filtering.

* Benefit in incremental/partial deployment: Any new solution SHOULD NOT assume pervasive adoption of the SAV method or the SAV-related information (e.g., Resource Public Key Infrastructure (RPKI) object registrations). It SHOULD benefit early adopters by providing effective protection from spoofing of source addresses even in partial deployment.

* Automatic updates to the SAV list and efficient convergence: Any new inter-domain SAV mechanism SHOULD be responsive to changes in the BGP (FIB/RIB) data, the SAV-related information (Section 2), or the SAV-specific information (Section 2). It SHOULD automatically update the SAV list while achieving efficient re-convergence of the same.

* Providing necessary security guarantee: If any proposed new SAV method requires exchanging SAV-related or SAV-specific information between ASes, security mechanisms SHOULD exist to assure trustworthiness of the information.

1.1. 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. Terminology

SAV List:
The table of prefixes that indicates the validity of a specific source IP address or source IP prefix per interface. Sometimes the terms 'RPF (Reverse Path Forwarding) list' or 'SAV rules' are used interchangeably with 'SAV list'.

Improper Block:
The validation results in packets with legitimate source addresses being blocked improperly due to an inaccurate SAV list.

Improper Permit:
The validation results in packets with spoofed source addresses being permitted improperly due to an inaccurate SAV list.

Customer Cone:
The Customer Cone (CC) of a given AS, denoted as AS-A, includes: (1) AS-A itself, (2) AS-A's direct customers (ASes), (3) The customers of AS-A's direct customers (indirect customers), (4) And so on, recursively, following all chains of provider-to-customer (P2C) links down the hierarchy.

Customer Cone Prefixes (CC Prefixes):
IP prefixes permitted by their owners to be originated by, or used as source addresses for data traffic originated from, one or more Autonomous Systems (ASes) within the CC.

SAV-related Information:
Objects registered using Resource Public Key Infrastructure (RPKI). This can include existing RPKI object types (e.g., ROAs and ASPAs) or any new type(s) that may be proposed.

SAV-specific Information:
Information dedicated to SAV, which may be defined and exchanged between ASes using a potentially new inter-AS communication protocol. The information may also be in the form of new RPKI object type(s) meant to assist SAV.

3. Existing Inter-domain SAV Mechanisms

Inter-domain SAV is typically performed at the AS level (on a per neighbor-AS-interface basis) and can be deployed at AS border routers (ASBRs) to prevent source address spoofing. There are various mechanisms available to implement inter-domain SAV for anti-spoofing ingress filtering [nist] [manrs] [isoc], which are reviewed in this section.

* ACL-based ingress filtering [RFC3704]: ACL-based ingress SAV filtering is a technique that relies on ACL rules to filter packets based on their source addresses. However, ACL-based ingress SAV filtering introduces significant operational overhead, as ACL rules need to be updated in a timely manner to reflect prefix or routing changes in the inter-domain routing system. One may think of using ACL as a disallow list on a provider interface to block source prefixes that are clearly invalid in the inter-domain routing context, such as IANA special purpose or unallocated IPv4/IPv6 prefixes, etc. But it is impractical to store and maintain a very large and dynamically varying set of unallocated IPv6 prefixes. Also, for the customer interfaces, the ACL method is impractical while other techniques (as described below) are more effective. ACL-based ingress SAV filtering has applicability for broadband cable or digital subscriber access loop (DSL) access networks where the service provider has clear knowledge of IP address prefixes it has allocated to manage those services. Here ACL can be used in an allow-list form.

* uRPF-based mechanisms: A class of SAV mechanisms are based on Unicast Reverse Path Forwarding (uRPF) [RFC3704] [RFC8704]. The core idea of uRPF for SAV is to exploit the symmetry of inter-domain routing: in many cases, the best next hop for a destination is also the best previous hop for the source. In other words, if a packet arrives from a certain interface, the source address of that packet should be reachable via the same interface, according to the FIB. However, symmetry in routing does not always hold in practice, and to address cases where it does not hold, many enhancements and modes of uRPF have evolved. Different modes of uRPF have different levels of strictness and flexibility, and network operators can choose from them to suit particular network scenarios. We briefly describe these modes as follows:

  - Strict uRPF [RFC3704]: Strict uRPF is the most stringent mode. It permits a packet only if it has a source address that is covered by a prefix in the FIB, and the next hop for that prefix is the same interface that the packet arrived on. This mode can be deployed at customer interfaces in some scenarios, e.g., a directly connected single-homed stub customer AS [nist].

  - Loose uRPF [RFC3704]: Loose uRPF verifies that the source address of a packet is routable on the internet by matching it with one or more prefixes in the FIB, regardless of the interface on which the packet arrives. If the source address is not routable, Loose uRPF discards the packet. Loose uRPF is typically deployed at the provider interfaces of an AS to block packets with source addresses from prefixes that are not routed on the global internet (e.g., IANA-allocated private-use addresses, unallocated IPv4/IPv6 addresses, multicast addresses, etc.).

- Feasible Path uRPF (FP-uRPF) [RFC3704]: Unlike Strict uRPF, which requires the packet to arrive on the exact best return path, FP-uRPF allows a packet to pass as long as the router could reach that source address through the interface it arrived on (based on the feasible routes in the Adj-RIBs-In [RFC4271]), even if the route isn’t the primary route (per best path selection). This makes it more effective in multi-homed environments where asymmetric routing is common, as it prevents legitimate traffic from being dropped simply because it didn’t take the "best" path back to the sender.

- Enhanced Feasible Path uRPF with Algorithm A (EFP-uRPF Alg-A) [RFC8704]: EFP-uRPF Alg-A expands the list of valid source addresses for a specific interface by including all prefixes associated with any Origin AS that is reachable through that interface. Instead of only accepting prefixes directly advertised on a link, the router identifies all the origin ASes present in the BGP updates received on that interface and then permits any prefix from those same ASes that it sees elsewhere in its Adj-RIBs-In (associated with all neighbors — customers, providers, peers). This "Origin AS-based" approach provides significantly more flexibility than strict or traditional FP-uRPF, as it accounts for cases where an AS in the CC may send traffic for one of its prefixes over a link where it only advertised a different prefix (multi-homing and asymmetric routing scenarios).

- Enhanced Feasible Path uRPF with Algorithm B (EFP-uRPF Alg-B) [RFC8704]: EFP-uRPF Alg-B provides even greater flexibility (compared to EFP-uRPF Alg-A) by aggregating all customer interfaces into a single "customer group" for validation purposes. The router first identifies all unique prefixes and origin ASes associated with all directly connected customer interfaces using only the Adj-RIBs-In associated with them. It then constructs a comprehensive RPF list that includes every prefix originated by those ASes, regardless of whether those prefixes were learned via customer, peer, or transit provider links. This list is applied uniformly across all customer-facing interfaces, attempting to ensure that legitimate traffic from a multihomed AS in the CC is never dropped, even if the traffic arrives on a different customer-facing port than the one where the specific prefix was advertised. In comparison to EFP-uRPF Alg-A, this method (Alg-B) reduces the possibility of improper block but at the expense of increased possibility of improper permit, i.e., reduced directionality.

- Virtual Routing and Forwarding (VRF) uRPF [RFC4364] [urpf] [manrs]: VRF uRPF uses a separate VRF table for each external BGP peer and is only a way of implementation for a SAV list.

4. Gap Analysis

Inter-domain SAV is essential for preventing source address spoofing traffic at all AS interfaces — customers, providers, and lateral peers. An ideal inter-domain SAV mechanism must block all spoofing traffic while permitting legitimate traffic in all scenarios of interest. However, in some cases, existing SAV mechanisms may unintentionally block legitimate traffic or permit spoofing traffic. This section aims to conduct a gap analysis of existing SAV mechanisms for different types of interfaces under various scenarios to identify their technical limitations.

4.1. SAV at Customer Interfaces

To prevent source address spoofing on customer interfaces, operators can enable ACL-based ingress filtering, or uRPF-based mechanisms such as Strict uRPF, FP-uRPF, or EFP-uRPF. However, the ACL method typically has high operational overhead. The uRPF-based mechanisms may cause improper block in two inter-domain scenarios: Limited Propagation of a Prefix (LPP) and Hidden Prefix (HP). They may also cause improper permit in the scenarios of source address Spoofing within a Customer Cone (SCC). The LPP scenario occurs when an AS applies traffic engineering (TE) using a no-export policy. One example is when an AS attaches NO_EXPORT BGP Community to some prefixes (routes) forwarded to some upstream providers (in multi-homing scenarios) (see Section 4.1.1). Sometimes this type of TE is done without attaching the NO_EXPORT, i.e., by selectively propagating different sets of prefixes to different upstream providers. The Hidden Prefix (HP) scenario is typically associated with the Direct Server Return (DSR) scenario; anycast prefix in a Content Delivery Network (CDN) application is not announced by the AS where the DSR (edge server) is located (see Section 4.1.2). Source address Spoofing within a Customer Cone (SCC) scenario arises when a prefix at one Autonomous System (AS) in the CC is spoofed from another AS in the same CC Section 4.1.3. It is recognized that unless there is full adoption of SAV in the customer cone (CC) of the interface in consideration, improper permit is not fully preventable in the SCC scenario.

Figure 1 provides an overview of the gaps associated with the ACL method, Strict uRPF, FP-uRPF, and EFP-uRPF for SAV at customer interfaces in the LPP, HP, and SCC scenarios mentioned above. Illustrations and analyses of these gaps are provided in Section 4.1.1, Section 4.1.2, and Section 4.1.3, respectively.

```
+--------------------+------------+-----------+-------+--------+
|Traffic & Scenarios |     ACL    |Strict uRPF|FP-uRPF|EFP-uRPF|
+----------+---------+------------+-----------+-------+--------+
|Legitimate|   LPP   |            |                            |
|Traffic   +---------+            |       Improper Block       |
|          |   HP    |    High    |         possible           |
+----------+---------+Operational-+--------------------+--------+
|          |         |  Overhead  |                    |Improper|
|Spoofed   |  no SCC |   (HOO)    |                    |Permit  |
|Traffic   |         |            |   Functions as     |only for|
|          |         |            |      Expected      |EFP-uRPF|
|          |         |            |                    |Alg-B   |
|+---------+---------+            +--------------------+--------|
|Spoofed   |   SCC   |            |                            |
|Traffic   |         |            |       Improper Permit      |
|          |         |            |   (in partial deployment)  |
+----------+---------+------------+----------------------------+
```

LPP = Limited Propagation of a Prefix
HP = Hidden Prefix
SCC = Spoofing within a CC
'Functions as Expected' connotes the absence of improper permit.
It also connotes low operational overhead.

Figure 1: The gaps of ACL-based ingress filtering, Strict uRPF, FP-uRPF, and EFP-uRPF for customer interfaces for the scenarios of interest.

4.1.1. Limited Propagation of a Prefix (LPP) Scenario

In inter-domain networks, some prefixes may not propagate from a customer to all its providers and/or may not propagate transitively from the providers to all their providers due to various factors, such as the use of NO_EXPORT or NO_ADVERTISE Communities, or some other selective-export policies meant for traffic engineering. In these cases, it is possible that a prefix (route) announcement in the CC associated with a customer interface has limited propagation in the CC and is not received on that interface. Then the prefix is invisible in BGP at that interface but the traffic with source address in that prefix may still be received on that interface. This can give rise to improper block when performing SAV with existing mechanisms. These mechanisms include EFP-uRPF Alg-A, which is the focus on in the following analysis, while it also applies to Strict uRPF and FP-uRPF. All these mechanisms suffer from the same problem of improper block in this scenario.

```
                             +-----------------+
                             |    AS 3(P3)     |
                             +-+/\-------+/\+--+
                                /         \
                               /           \
                              /             \
                             / (C2P)         \
                    +-------------------+     \
                    |     AS 4(P4)      |      \
                    ++/\+--+/\+-----+/\++       \
                      /     |        \           \
            P2[AS 2] /      |         \           \
                    /       |          \           \
                   / (C2P)  |           \ P5[AS 5]  \ P5[AS 5]
   +-----------------+      |            \           \
   |    AS 2(P2)     |      | P1[AS 1]    \           \
   +----------+/\+--+       | P6[AS 1]     \           \
                \           |               \           \
        P1[AS 1] \          |                \           \
        NO_EXPORT \         |                 \           \
                   \(C2P)   |(C2P)        (C2P)\      (C2P)\
                 +-----------------+         +-----------------+
                 |  AS 1(P1, P6)   |         |    AS 5(P5)     |
                 +-----------------+         +-----------------+
```

Figure 2: Limited propagation of a prefix caused by NO_EXPORT.

In the scenario of Figure 2, AS 1 is a customer of AS 2; AS 1 and AS 2 are customers of AS 4; AS 4 is a customer of AS 3; and AS 5 is a customer of both AS 3 and AS 4. AS 1 advertises prefixes P1 to AS 2 with the NO_EXPORT community attribute attached, preventing AS 2 from further propagating the route for prefix P1 to AS 4. Consequently, AS 4 only learns the route for prefix P1 from AS 1 in this scenario. Suppose AS 1 and AS 4 have deployed inter-domain SAV while other ASes have not, and AS 4 has deployed EFP-uRPF at its customer interfaces.

If AS 4 deploys EFP-uRPF Alg-A at customer interfaces, it will require packets with source addresses in P1 or P6 to only arrive on the interface with AS 1. When AS 1 sends legitimate packets with source addresses in P1 or P6 to AS 4 via AS 2, AS 4 improperly blocks these packets. The same improper block problem occurs with the use of Strict uRPF or FP-uRPF. EFP-uRPF with Alg-B can avoid the improper block in this specific scenario, but even this SAV method would have the improper block if the TE at AS 1 is such that none of the customer interfaces at AS 4 receives a route for P1 (or P6).

4.1.2. Hidden Prefix (HP) Scenario

CDNs use the concepts of anycast [RFC4786][RFC7094] and DSR to improve the quality of service by placing edge servers with content closer to users. An anycast IP address is assigned to devices in different locations, and incoming requests are routed to the closest edge server (DSR) location. Usually, only locations with rich connectivity announce the anycast IP address through BGP. The CDN server receives requests from users and creates tunnels to the edge locations, from where content is sent directly to users. DSR requires servers in the edge locations to use the anycast IP address as the source address in response packets. However, the ASes serving the edge servers do not announce the anycast prefixes through BGP, so the anycast prefix is hidden (invisible in BGP) on the customer interface side at intermediate ASes which — with existing inter-domain SAV mechanisms — would improperly block the response packets.

Figure 3 illustrates a DSR scenario where the anycast IP prefix P3 is advertised by AS 3 through BGP. In this example, AS 3 is the provider of AS 4 and AS 5; AS 4 is the provider of AS 1, AS 2, and AS 5; and AS 2 is the provider of AS 1. AS 2 and AS 4 have deployed inter-domain SAV. When a user at AS 2 sends a request to the anycast destination IP, the forwarding path is AS 2->AS 4->AS 3. The anycast server in AS 3 receives the request and tunnels it to the edge servers in AS 1. Finally, the edge server sends the content packets to the user with source addresses in prefix P3. Let us say, the forwarding path for the content packets is AS 1-> AS 4->AS 2. Since AS 4 does not receive routing information for prefix P3 from AS 1, EFP-uRPF Alg-A or EFP-uRPF Alg-B (or any other existing uRPF-based mechanism) at the customer interface of AS 4 facing AS 1 will improperly block the response packets from AS 1.

```
                                   +----------------+
                   Anycast Server+-+    AS 3(P3)    |
                                   +-+/\----+/\+----+
                                      /       \
                            P3[AS 3] /         \
                                    /           \
                                   / (C2P)       \
                          +----------------+      \
                          |    AS 4(P4)    |       \
                          ++/\+--+/\+--+/\++        \
             P6[AS 2, AS 1] /     |      \           \
            P1[AS 2, AS 1] /      |       \           \
                 P2[AS 2] /       |        \           \
                         / (C2P)  |         \ P5[AS 5]  \ P5[AS 5]
         +----------------+       |          \           \
   User+-+    AS 2(P2)    |       | P1[AS 1]  \           \
         +----------+/\+--+       | P6[AS 1]   \           \
                      \           |             \           \
              P6[AS 1] \          |              \           \
               P1[AS 1] \         |               \           \
                         \(C2P)   |(C2P)      (C2P)\      (C2P)\
                       +---------------+         +----------------+
          Edge Server+-+  AS 1(P1, P6)  |        |    AS 5(P5)    |
                       +----------------+        +----------------+
```

P3 is the anycast prefix and is only advertised by AS 3 through BGP.

Figure 3: A Direct Server Return (DSR) scenario.

Further, there are cases of specific prefixes that may be exclusively used as source addresses (legitimately) without being advertised via BGP by any AS. While different from DSR scenarios, these cases similarly result in existing inter-domain SAV mechanisms improperly blocking legitimate traffic originating from such prefixes.

4.1.3. Source Address Spoofing within a Customer Cone (SCC) Scenario

In general, improper permit of spoofed packets in SCC scenarios is unavoidable for various uRPF-based methods in partial deployment. For example, consider a topology in which AS 1 and AS 2 are customers of AS 3; and AS 3 is a customer of AS 4. AS 1 and AS 2 originate prefixes P1 and P2, respectively. AS 4 performs SAV on its customer interface with AS 3. P1 and P2 are announced from AS 3 to AS 4 and they would be included in the SAV list (allowlist) of AS 4 with any SAV mechanism. Assume AS 3 doesn’t do SAV. Now as an example of SCC, if AS 2 spoofs AS 1’s prefix P1 and sends the spoofed packets to AS 4 (via AS 3), there is no way for AS 4 to detect the spoofed traffic. AS 4’s SAV cannot differentiate between the spoofed and the legitimate packets that have source address in P1. In an SCC scenario of this nature, the only recourse for blocking the spoofed traffic is for AS 3 also to be upgraded to do SAV, i.e., deployment of SAV closer to the source of spoofing.

Another scenario is highlighted in Figure 4 while using EFP-uRPF Alg-B method on customer interfaces. This scenario is non-SCC from the perspective of each individual customer interfaces of AS 4, but it is SCC from the perspective of AS 4 as a whole. EFP-uRPF Alg-B relaxes directionality to reduce (or eliminate) false positives and that makes it more susceptible to SCC (per the latter perspective). This is expected because EFP-uRPF Alg-B somewhat conservatively applies the same relaxed SAV list across all customer interfaces.

```
                                        +----------------+
                                        |    AS 3(P3)    |
                                        +-+/\----+/\+----+
                                           /       \
                                          /         \
                                         /           \
                                        / (C2P)       \
                               +----------------+      \
                               |    AS 4(P4)    |       \
                               ++/\+--+/\+--+/\++        \
                  P6[AS 2, AS 1] /     |      \           \
                 P1[AS 2, AS 1] /      |       \           \
                      P2[AS 2] /       |        \           \
                              / (C2P)  |         \ P5[AS 5]  \ P5[AS 5]
              +----------------+       |          \           \
 Spoofer(P5’)-+    AS 2(P2)    |       | P1[AS 1]  \           \
              +----------+/\+--+       | P6[AS 1]   \           \
                           \           |             \           \
                   P6[AS 1] \          |              \           \
                    P1[AS 1] \         |               \           \
                              \(C2P)   |(C2P)      (C2P)\      (C2P)\
                            +----------------+        +----------------+
                            |  AS 1(P1, P6)  |        |    AS 5(P5)    |
                            +----------------+        +----------------+
 P5’ is the spoofed source prefix P5 by the spoofer which is inside of
 AS 2 or connected to AS 2 through other ASes.
```

Figure 4: A scenario of source address spoofing within a customer cone.

In Figure 4, the source address spoofing takes place within AS 4’s customer cone, where the spoofer, which is inside of AS 2 or connected to AS 2 through other ASes, sends spoofing traffic with spoofed source addresses in P5 to AS 3 along the path AS 2->AS 4-> AS 3. The arrows in Figure 4 illustrate the commercial relationships between ASes. AS 3 serves as the provider for AS 4 and AS 5, while AS 4 acts as the provider for AS 1, AS 2, and AS 5. Additionally, AS 2 is the provider for AS 1. Suppose AS 1 and AS 4 have deployed inter-domain SAV, while the other ASes have not.

If AS 4 deploys EFP-uRPF Alg-B at its customer interfaces, it will allow packets with source addresses in P5 to originate from AS 1, AS 2, and AS 5. Consequently, AS 4 will improperly permit the spoofed packets from AS 2, enabling them to propagate further.

In the scenario of Figure 4, Strict uRPF, FP-uRPF, and EFP-uRPF Alg-A — applied on the customer interfaces — work effectively to block the spoofed packets from AS 2. This is because these mechanisms have stronger directionality property than EFP-uRPF Alg-B.

4.2. SAV at Peer Interfaces

SAV is used at peer interfaces for validating the traffic entering the validating AS and destined for the AS’s customer cone. The data packets received from a customer or lateral peer AS must have source addresses belonging only to the prefixes in the customer cone (CC) of that AS. In both cases, the focus is on discovering all prefixes in the CC of the neighbor AS. So, in principle, the SAV techniques suitable on a customer interface may also be used on a peer interface, especially EFP-uRPF Alg-A or Alg-B, which are more accommodative of asymmetric routing. Indeed, asymmetric routing is thought to be prevalent for peer interfaces. If SAV techniques suitable for customer interfaces are considered for peer interfaces, then the gap analysis of Section 4.1 would also be applicable to the SAV for the peer interfaces. However, due to increased concern about asymmetric routing, network operators may conservatively use the same relaxed SAV techniques for peer interfaces as those for provider interfaces, e.g., Loose uRPF Section 4.3. In that case, the gap analysis of Section 4.3 would also be applicable to the SAV for peer interfaces.

4.3. SAV at Provider Interfaces

SAV is used at provider interfaces for validating the traffic entering the AS and destined for the AS’s customer cone. Figure 5 summarizes the gaps of ACL-based ingress filtering and Loose uRPF for SAV at provider interfaces in the scenarios of interest. ACL-based ingress filtering may effectively block spoofing traffic from provider AS, while appropriately allowing legitimate traffic, but it has high operational overhead. On the other hand, Loose uRPF correctly permits legitimate traffic, but it can also mistakenly allow spoofing traffic to pass through.

In Figure 5, Spoofed from Provider Tree (SPT) is a scenario where the spoofed traffic comes from the provider tree, i.e., the providers in the transitive hierarchy above the validating AS. The spoofed prefix may belong to (originated by) any AS in the Internet other than the spoofing AS; it may even belong to an AS in the customer cone of the validating AS (example below).

Traffic & Scenarios		ACL	Loose uRPF
Legitimate Traffic	--	High Operational Overhead (HOO)	Functions as Expected
Spoofed Traffic	Spoofed from Provider Tree (SPT)		Improper Permit

'Functions as Expected' connotes the absence of improper block. It also connotes low operational overhead.

Figure 5: The gaps of ACL-based ingress filtering and Loose uRPF at provider interfaces in the scenarios of interest.

Figure 6 illustrates a scenario of SPT and is used to analyze the gaps of ACL-based ingress filtering and Loose uRPF.

```
                             +----------------+
               Spoofer(P1’)+-+    AS 3(P3)    |
                             +-+/\----+/\+----+
                                /       \
                               /         \
                              /           \
                             / (C2P)       \
                    +----------------+      \
                    |    AS 4(P4)    |       \
                    ++/\+--+/\+--+/\++        \
       P6[AS 2, AS 1] /     |      \           \
      P1[AS 2, AS 1] /      |       \           \
           P2[AS 2] /       |        \           \
                   / (C2P)  |         \ P5[AS 5]  \ P5[AS 5]
   +----------------+       |          \           \
   |    AS 2(P2)    |       | P1[AS 1]  \           \
   +----------+/\+--+       | P6[AS 1]   \           \
                \           |             \           \
        P6[AS 1] \          |              \           \
         P1[AS 1] \         |               \           \
                   \ (C2P)  | (C2P)    (C2P) \     (C2P) \
                  +----------------+        +----------------+
                  |  AS 1(P1, P6)  |        |    AS 5(P5)    |
                  +----------------+        +----------------+
```

P1’ is the spoofed source prefix P1 by the spoofer which is inside of AS 3 or connected to AS 3 through other ASes.

Figure 6: A scenario of source address spoofing from provider AS.

In Figure 6, the spoofer which is inside of AS 3 or connected to AS 3 through other ASes forges the source addresses in P1 and sends the spoofing traffic to the destination addresses in P2 at AS 2. AS 1 is a customer of AS 2; AS 1 and AS 2 are customers of AS 4; AS 4 is a customer of AS 3; and AS 5 is a customer of both AS 3 and AS 4. Suppose AS 4 and AS 1 have deployed inter-domain SAV, while the other ASes have not.

Using the ACL method in the form of a disallow (deny) list at the provider interface of AS 4 (facing AS 3) incurs a very high operational overhead. As mentioned before (Section 3), it is impractical to store and maintain a very large and dynamically varying set of unallocated IPv6 prefixes in the ACL.

Applying Loose uRPF at the provider interface of AS 4 (facing AS 3) can greatly reduce the operational overhead because it uses the FIB as the information source for allowed prefixes, and can adapt to changes in the network to prevent false positives (improper blocking). However, using Loose uRPF at AS 4 will naturally permit packets with source addresses in P1 (since P1 is present in the FIB) and hence will not prevent the improper permit of the spoofed packets from AS 3 Figure 6. This is an expected limitation of Loose uRPF.

5. Problem Statement

Figure 7 provides a comprehensive summary of the gap analysis in Section 4. It highlights the scenarios where existing inter-domain SAV mechanisms may encounter issues, including instances of improper blocking of legitimate traffic, improper permitting of spoofing traffic, or high operational overhead. The various entries in the table in Section 4 can be traced back to the terminology and analyses presented in Section 4.

Problems	ACL (CI or PI)	Strict uRPF (CI)	Loose uRPF (PI)	FP-uRPF (CI)	EFP-uRPF (CI)
Improper Block	YES/NO (manual operator diligence)	YES (LPP, HP)	NO**	YES (LPP, HP)	
Improper Permit	YES/NO (manual operator diligence)	NO (no SCC) YES (SCC)	YES (SPT)	NO (no SCC) YES (SCC)	
HOO	YES (Any Scenarios)	NO			

CI = Customer Interface
PI = Provider Interface
HOO = High Operational Overhead
LPP = Limited Propagation of a Prefix
HP = Hidden Prefix
SCC = Spoofing within a CC
SPT = Spoofing from Provider Tree
** Typically, an HP (like DSR prefixes) is hidden on the CIs but received on a provider or peer interface; hence included in the FIB and that helps avoid improper block for Loose uRPF.

Figure 7: The scenarios where existing inter-domain SAV mechanisms may have improper block problem for legitimate traffic, improper permit problem for spoofing traffic, or high operational overhead.

The problem statement that results from the gap analysis can be expressed as follows. New proposals for SAV should aim to fill in the following problem areas (gaps) found in the currently standardized SAV methods (found in IETF RFCs):

* Improper block: Existing uRPF-based mechanisms suffer from improper block (false positives) in two inter-domain scenarios: limited propagation of a prefix (e.g., NO_EXPORT and some other traffic engineering (TE) scenarios) and hidden prefix (e.g., CDN/DSR scenario).

* Improper permit: With some existing uRPF-based SAV mechanisms, improper permit (false negatives) can happen on any type of interface (customer, lateral peer, or provider). Specifically, if the method relaxes the directionality constraint [RFC3704] [RFC8704]} to try to achieve zero improper blocking, the possibility of improper permit increases. (Note: It is recognized that unless there is full adoption of SAV in the customer cone (CC) of the interface in consideration, improper permit is not fully preventable in scenarios where source address spoofing occurs from within the CC, i.e., a prefix at one Autonomous System (AS) in the CC is spoofed from another AS in the same CC.)

* High operational overhead (HOO): ACL-based ingress SAV filtering introduces significant operational overhead, as it needs to update ACL rules manually to adapt to prefix or routing changes in a timely manner. The HOO issue does not pertain to existing uRPF-based mechanisms.

The limitations of existing uRPF-based mechanisms are due to their exclusive reliance on BGP data. Although the algorithms themselves have evolved (e.g., [RFC8704]), the underlying input has remained unchanged, inherently constraining their accuracy in scenarios such as LPP and HP. With the availability of authoritative SAV-related information, plus the potential SAV-specific information (Section 4), it would be possible to develop comprehensive new SAV algorithms or mechanisms to overcome the existing gaps.

6. Requirements for New Inter-domain SAV Mechanisms

This section lists the requirements for any new inter-domain SAV mechanisms which may be proposed to bridge the technical gaps of existing mechanisms.

6.1. Accurate Validation

Any new inter-domain SAV mechanism MUST avoid improper blocking and have superior directionality property (reject more spoofed traffic) than existing inter-domain SAV mechanisms. The requirement applies for all directions of AS peering (customer, provider, and peer).

6.2. Reducing Operational Overhead

Any new inter-domain SAV mechanism MUST be able to automatically adapt to network dynamics and asymmetric routing scenarios. Any such solution MUST have less operational overhead than ACL-based ingress SAV filtering.

6.3. Early Adopters Benefit in Incremental/Partial Deployment

Any new solution SHOULD NOT assume pervasive adoption of the SAV method or the SAV-related information (e.g., Resource Public Key Infrastructure (RPKI) objects such as ROAs and ASPAs). It SHOULD benefit early adopters by providing effective protection from spoofing of source addresses even in partial deployment.

6.4. Providing Necessary Security Guarantee

SAV-related information, such as RPKI objects, may be used for designing a more accurate SAV. Such information must be protected at their repositories and during communication to the relying parties (the BGP security community is already diligent about this). If any proposed SAV method requires exchanging SAV-specific information between ASes, security mechanisms must exist to assure trustworthiness of the information. The idea is to prevent malicious injection or alteration of the SAV-specific information.

6.5. Automatic Updates to the SAV List and Efficient Convergence

Any new inter-domain SAV mechanism SHOULD be responsive to changes in the BGP (FIB/RIB) data, the SAV-related information (Section 2), or the SAV-specific information (Section 2). It SHOULD automatically update the SAV list while achieving efficient re-convergence of the same. In this context, convergence refers to the stabilization of the SAV lists on the AS-to-AS interfaces performing SAV. It is essential that any new SAV mechanism converges to the correct updated SAV list in a proper manner, minimizing both improper block and improper permit during the process.

7. Inter-domain SAV Scope

Any new inter-domain SAV mechanisms should work in the same Internet Protocol (IP) address scenarios as existing SAV methods do. Generally, it includes all IP-encapsulated scenarios:

* Native IP forwarding: This includes both the global routing table based forwarding and Customer Edge (CE) site forwarding of VPN traffic.

* IP-encapsulated Tunnel (IPsec, GRE, SRv6, etc.): In this scenario, the focus is on the validation of the outer layer IP source address.

* Both IPv4 and IPv6 addresses.

The scope does not include:

* Non-IP packets: This includes MPLS label-based forwarding and other non-IP-based forwarding.

In addition, any new inter-domain SAV mechanisms MUST NOT modify the data plane packets. Existing architectures or protocols or mechanisms can be inherited by any such mechanism to achieve better SAV effectiveness.

8. Security Considerations

The SAV list will be generated based on routing information from BGP (FIB/RIB), SAV-related information, and/or SAV-specific information. If the information is poisoned by attackers, the SAV list will be inaccurate. Legitimate packets may be dropped improperly or malicious traffic with spoofed source addresses may be permitted improperly. BGP routing security using available methods for the prevention, detection, and mitigation of route hijacks, route leaks, and AS_PATH manipulations should be deployed which leads to greater accuracy of the BGP (FIB/RIB) information used for computing SAV lists.

9. IANA Considerations

This document does not request any IANA allocations.

10. Contributors

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

11. References

11.1. 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/rfc/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/rfc/rfc8174>.

11.2. Informative References

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/rfc/rfc4271>.

[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/rfc/rfc5210>.

[RFC7908] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, DOI 10.17487/RFC7908, June 2016, <https://www.rfc-editor.org/rfc/rfc7908>.

[RFC9234] Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages", RFC 9234, DOI 10.17487/RFC9234, May 2022, <https://www.rfc-editor.org/rfc/rfc9234>.

[RFC7094] McPherson, D., Oran, D., Thaler, D., and E. Osterweil, "Architectural Considerations of IP Anycast", RFC 7094, DOI 10.17487/RFC7094, January 2014, <https://www.rfc-editor.org/rfc/rfc7094>.

[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/rfc/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/rfc/rfc8704>.

[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/rfc/rfc2827>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <https://www.rfc-editor.org/rfc/rfc4364>.

[RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, DOI 10.17487/RFC4786, December 2006, <https://www.rfc-editor.org/rfc/rfc4786>.

[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-21, 18 January 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-21>.

[manrs] MANRS, "Anti-Spoofing - Preventing traffic with spoofed source IP addresses (Module 5)", <https://manrs.org/resources/training/tutorials/anti-spoofing/>.

[isoc] Internet Society, "Addressing the challenge of IP spoofing", 2015, <https://www.internetsociety.org/resources/doc/2015/addressing-the-challenge-of-ip-spoofing/>.

[nist] Sriram, K. and D. Montgomery, "Border Gateway Protocol Security and Resilience", NIST SP 800-189r1 , 2025, <https://doi.org/10.6028/NIST.SP.800-189r1.ipd>.

[urpf]   Cisco Systems, Inc., "Unicast Reverse Path Forwarding Enhancements for the Internet Service Provider-Internet Service Provider Network Edge", 2005, <https://www.cisco.com/c/dam/en_us/about/security/intelligence/urpf.pdf>.

Acknowledgements

Many thanks to Jared Mauch, Barry Greene, Fang Gao, Anthony Somerset, Yuanyuan Zhang, Igor Lubashev, Alvaro Retana, Joel Halpern, Ron Bonica, Aijun Wang, Michael Richardson, Li Chen, Gert Doering, Mingxing Liu, John O'Brien, and Roland Dobbins for their reviews, comments, and suggestions. Apologies to any others whose names the authors may have inadvertently missed mentioning.

Authors' Addresses

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

Lancheng Qin
Zhongguancun Laboratory
Beijing
China
Email: qinlc@zgclab.edu.cn

Libin Liu
Zhongguancun Laboratory
Beijing
China
Email: liulb@zgclab.edu.cn

Mingqing Huang
Huawei
Beijing
China
Email: huangmingqing@huawei.com

Kotikalapudi Sriram
USA National Institute of Standards and Technology
Gaithersburg, MD
United States of America
Email: sriram.ietf@gmail.com