Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 30 January 2026

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

29 July 2025

Source Address Validation in Inter-domain Networks Gap Analysis, Problem Statement, and Requirements
draft-ietf-savnet-inter-domain-problem-statement-10

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

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 30 January 2026.

Copyright Notice

Table of Contents

1. Introduction 2
1.1. Requirements Language 4
2. Terminology 4
3. Existing Inter-domain SAV Mechanisms 5
4. Gap Analysis 7
4.1. SAV at Customer Interfaces 7
4.1.1. Limited Propagation of Prefixes 9
4.1.2. Hidden Prefixes 10
4.1.3. Source Address Spoofing within a Customer Cone . . . 12
4.2. SAV at Provider/Peer Interfaces 14
4.2.1. Source Address Spoofing from Provider/Peer AS 15
5. Problem Statement 16
6. Requirements for New Inter-domain SAV Mechanisms 18
6.1. Accurate Validation 18
6.1.1. Improving Validation Accuracy over Existing Mechanisms 19
6.1.2. Working in Incremental/Partial Deployment 20
6.1.3. Providing Necessary Security Guarantee 20
6.2. Automatic Update 20
6.2.1. Reducing Operational Overhead 20
6.2.2. Guaranteeing Convergence 20
7. Inter-domain SAV Scope 21
8. Security Considerations 21
9. IANA Considerations 21
10. Contributors 22
11. References 22
11.1. Normative References 22
11.2. Informative References 23
Acknowledgements 23
Authors' Addresses 23

1. Introduction

Source address validation (SAV) is crucial for protecting networks from source address (SA) spoofing attacks [RFC2827] [RFC3704] [RFC8704]. The MANRS initiative advocates deploying SAV as close to the source as possible [manrs], and access networks are the first line of defense against source address spoofing. However, access networks face various challenges in deploying SAV mechanisms due to different network environments, router vendors, and operational preferences. Hence, SAV may not be deployed ubiquitously in access

networks. In addition, SA spoofing may also originate in ISP networks at higher levels of heirarchy in the Internet. So, deployment of SAV mechanisms in the edge routers of enterprises as well as the ISP networks (at different heirarchal levels or tiers) is needed to prevent source address spoofing along the data forwarding paths. [RFC5210] highlighted the importance of SAV at various network locations: access, intra-domain, and inter-domain. This document focuses on providing gap analysis and describing the problem space of existing inter-domain SAV solutions, and defining the requirements for a new solution of these problems. Access Control List (ACL) and unicast Reverse Path Forwarding (uRPF) techniques are currently utilized for inter-domain SAV [RFC3704] [RFC8704]. Here only existing IETF RFCs are considered as the state of the art (BCP 38 [RFC2827] and BCP 84 [RFC3704] [RFC8704]); IETF works-in-progress are not included in that.

There are several existing mechanisms for inter-domain SAV. This document analyzes them and attempts to answer: i) what are the technical gaps (Section 4), ii) what are the fundamental problems (Section 5), and iii) what are the practical requirements for the solution of 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 in two inter-domain scenarios: limited propagation of prefixes and hidden prefixes.

* Improper permit: Existing uRPF-based mechanisms exhibit improper permit in scenarios involving source address spoofing within a customer cone or from a provider/peer AS.

* High operational overhead: 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.

To address these problems, in Section 6, this document outlines the following technical requirements for a new solution:

* Improving validation accuracy over existing mechanisms: A new solution MUST avoid improper block and minimize improper permit.

* Reducing operational overhead: A new solution MUST have less operational overhead than ACL-based ingress SAV filtering.

In addition, this document defines three more requirements to ensure practicality:

* Working in incremental/partial deployment: A new solution MUST NOT assume pervasive adoption including the adoption of both SAV and SAV-related information and SHOULD provide effective protection for source addresses when it is partially deployed in the Internet.

* Providing necessary security guarantee: A new solution SHOULD secure the communicated information between ASes if it requires exchanging specific information between ASes.

* Guaranteeing convergence: A new solution SHOULD achieve accurate SAV rule convergence in response to prefix or routing changes.

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 Rule:
: The rule that indicates the validity of a specific source IP address or source IP prefix.

Improper Block:
: The validation results that the packets with legitimate source addresses are blocked improperly due to inaccurate SAV rules.

Improper Permit:
: The validation results that the packets with spoofed source addresses are permitted improperly due to inaccurate SAV rules.

Active forwarding paths:
: The paths that the legitimate traffic goes through in the data plane at a given time period.

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. It can be applied at provider interfaces, peer interfaces, or customer interfaces of an AS, and is recommended for deployment at provider interfaces [manrs]. At the provider interfaces, ACL-based ingress SAV filtering can block source prefixes that are clearly invalid in the inter-domain routing context, such as IANA special purpose or unallocated IPv4/IPv6 prefixes and the AS's internal-only prefixes. 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. It is also impractical to store a very large and dynamically varying unallocated IPv6 prefixes. At the customer interfaces, ACL-based ingress filtering is less desirable. Other techniques (as described below) are more effective for ingress SAV filtering on customer interfaces. ACL-based ingress SAV filtering has applicability for broadband cable or digital subscriber access loop (DSL) access networks where the service provider has clear knwoledge of IP address prefixes it has allocated to manage those services.

* uRPF-based mechanisms: A class of SAV mechanisms are based on Unicast Reverse Path Forwarding (uRPF) [RFC3704]. 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 holds in practice, and to address cases where it does not hold, many enhancements and modes of uRPF are proposed. Different modes of uRPF have different levels of strictness and flexibility, and network operators can choose from them to suit particular network scenarios. We describe these modes as follows:

  - Strict uRPF [RFC3704]: Strict uRPF is the most stringent mode, and it only permits packets that have a source address that is covered by a prefix in the FIB, and that the next hop for that

prefix is the same as the incoming interface. This mode is recommended for deployment at customer interfaces that directly connect to an AS with suballocated address space, as it can prevent spoofing attacks from that AS or its downstream ASes [nist].

- Loose uRPF [RFC3704]: Loose uRPF verifies that the source address of the packet is routable in the Internet by matching it with one or more prefixes in the FIB, regardless of which interface the packet arrives at. 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 that are obviously disallowed, such as non-global prefixes (e.g., private addresses, multicast addresses, etc.) or the prefixes that belong to the customer AS itself [nist].

- Feasible Path uRPF (FP-uRPF) [RFC3704]: maintains a reverse path forwarding (RPF) list, which contains the prefixes and all their permissible routes including the optimal and alternative ones. It permits an incoming packet only if the packet’s source address is encompassed in the prefixes of the RPF list and its incoming interface is included in the permissible routes of the corresponding prefix. FP-uRPF is recommended to be deployed at customer interfaces or peer interfaces, especially those that are connected to multi-homed customer ASes [nist].

- 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 table.

- Enhanced Feasible Path uRPF (EFP-uRPF) [RFC8704]: EFP-uRPF is based on the principle that if BGP updates for multiple prefixes with the same origin AS were received on different interfaces (at border routers), then incoming data packets with source addresses in any of those prefixes should be accepted on any of those interfaces. The EFP-uRPF specification provides two alternate algorithms: Algorithm A which is stricter with a greater sense of directionality and Algorithm B which is more permissive with a lesser sense of directionality. EFP-uRPF can more effectively accomodate asymmetric routing scenarios than FP-uRPF. EFP-uRPF is a part of BCP84. EFP-uRPF can be used at customer as well as lateral peer interfaces of an AS. It is not deployed yet in the Internet.

* Carrier Grade NAT (CGN): CGN is a network technology used by service providers to translate between private and public IPv4 addresses within their network. CGN enables service providers to assign private IPv4 addresses to their customer ASes instead of public, globally unique IPv4 addresses. The private side of the CGN faces the customer ASes, and when an incoming packet is received from a customer AS, CGN checks its source address. If the source address is included in the address list of the CGN's private side, CGN performs address translation. Otherwise, it forwards the packet without translation. However, since CGN cannot determine whether the source address of an incoming packet is spoofed or not, additional SAV mechanisms need to be implemented to prevent source address spoofing [manrs].

* BGP origin validation (BGP-OV) [RFC6811]: Attackers can bypass uRPF-based SAV mechanisms by using prefix hijacking in combination with source address spoofing. By announcing a less-specific prefix that does not have a legitimate announcement, the attacker can deceive existing uRPF-based SAV mechanisms and successfully perform address spoofing. To protect against this type of attack, a combination of BGP-OV and uRPF-based mechanisms like FP-uRPF or EFP-uRPF is recommended [nist]. BGP routers can use ROA information, which is a validated list of {prefix, maximum length, origin AS}, to mitigate the risk of prefix hijacks in advertised routes.

4. Gap Analysis

Inter-domain SAV is essential in preventing source address spoofing traffic across all AS interfaces, including those of customers, providers, and peers. An ideal inter-domain SAV mechanism MUST block all spoofing traffic while permitting legitimate traffic in all scenarios. 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 used in the corresponding interfaces of these scenarios to identify their technical limitations.

4.1. SAV at Customer Interfaces

SAV is used at customer interfaces to validate traffic from the customer cone, including both legitimate traffic and spoofing traffic. To prevent the source address spoofing, operators can enable ACL-based ingress filtering and/or uRPF-based mechanisms at customer interfaces, namely Strict uRPF, FP-uRPF, or EFP-uRPF. However, uRPF-based mechanisms may cause improper block problems in two inter-domain scenarios: limited propagation of prefixes and hidden prefixes, or may cause improper permit problems in the

scenarios of source address spoofing within a customer cone, while ACL-based SAV ingress filtering needs to update SAV rules in a timely manner and lead to high operational overhead.

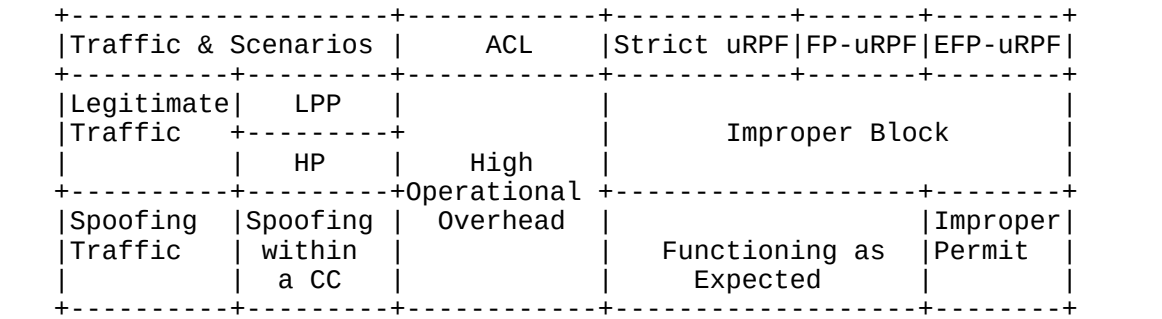

"LPP" represents a class of scenario called limited propagation of prefixes.
"HP" represents a class of scenario called hidden prefixes.
"Spoofing within a CC" represents a class of scenario where spoofing traffic occurs within a customer cone (CC) and the spoofed source addresses belong to this customer cone.
"Functioning as Expected" represents the inter-domain SAV mechanism does not cause improper block for legitimate traffic or improper permit for spoofing traffic in the corresponding scenarios, and has low operational overhead.

Figure 1: The gaps of ACL-based ingress filtering, Strict uRPF, FP-uRPF, and EFP-uRPF in the corresponding scenarios.

Figure 1 provides an overview of the gaps associated with ACL-based ingress filtering, Strict uRPF, FP-uRPF, and EFP-uRPF for SAV at customer interfaces in the corresponding scenarios. ACL-based ingress filtering has high operational overhead as performing SAV at customer interfaces. Strict uRPF, FP-uRPF, and EFP-uRPF, on the other hand, may incorrectly block legitimate traffic in the scenarios of limited propagation of prefixes or hidden prefixes. Furthermore, in the scenarios of source address spoofing within a customer cone, EFP-uRPF with algorithm B may inadvertently permit the spoofing traffic.

In the following, we analyze the gaps of Strict uRPF, FP-uRPF, and EFP-uRPF for SAV at customer interfaces in scenarios of limited propagation of prefixes, hidden prefixes, and source address spoofing within a customer cone, respectively.

4.1.1. Limited Propagation of Prefixes

In inter-domain networks, some prefixes may not be propagated to all domains due to various factors, such as NO_EXPORT or NO_ADVERTISE communities or other route filtering policies. This may cause asymmetric routing in the inter-domain context, which may lead to improper block when performing SAV with existing mechanisms. These mechanisms include EFP-uRPF, which we focus on in the following analysis, as well as Strict uRPF and FP-uRPF. All these mechanisms suffer from the same problem of improper block in this scenario.

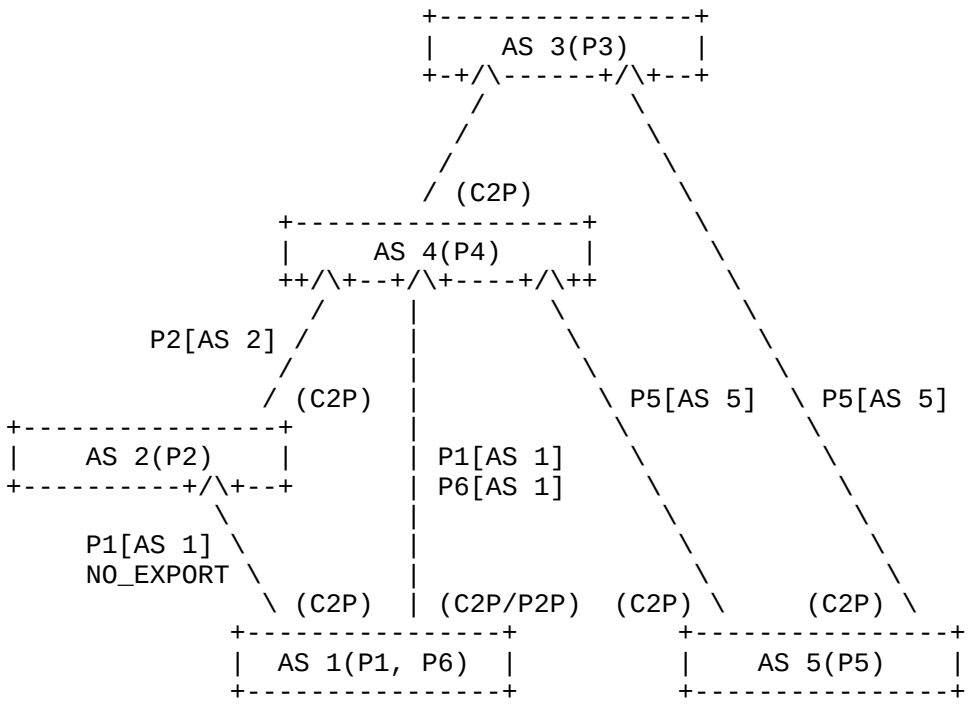

Figure 2: Limited propagation of prefixes caused by NO_EXPORT.

Figure 2 presents a scenario where the limited propagation of prefixes occurs due to the NO_EXPORT community attribute. In this scenario, AS 1 is a customer of AS 2, AS 2 is a customer of AS 4, AS 4 is a customer of AS 3, and AS 5 is a customer of both AS 3 and AS 4. The relationship between AS 1 and AS 4 can be either customer-to-provider (C2P) or peer-to-peer (P2P). AS 1 advertises prefixes P1 to AS 2 and adds the NO_EXPORT community attribute to the BGP advertisement sent to AS 2, 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.

Assuming that AS 1 is the customer of AS 4, if AS 4 deploys EFP-uRPF with algorithm A at customer interfaces, it will require packets with source addresses in P1 or P6 to only arrive from AS 1. When AS 1 sends legitimate packets with source addresses in P1 or P6 to AS 4 through AS 2, AS 4 improperly blocks these packets. The same problem applies to Strict uRPF and FP-uRPF. Although EFP-uRPF with algorithm B can avoid improper block in this case, network operators need to first determine whether limited prefix propagation exists before choosing the suitable EFP-uRPF algorithms, which adds more complexity and overhead to network operators. Furthermore, EFP-uRPF with algorithm B is not without its problems. For example, if AS 1 is the peer of AS 4, AS 4 will not learn the route of P1 and P6 from its customer interfaces. In such case, both EFP-uRPF with algorithm A and algorithm B have improper block problems.

4.1.2. Hidden Prefixes

Some servers’ source addresses are not advertised through BGP to other ASes. These addresses are unknown to the inter-domain routing system and are called hidden prefixes. Legitimate traffic with these hidden prefixes may be dropped by existing inter-domain SAV mechanisms, such as Strict uRPF, FP-uRPF, or EFP-uRPF, because they do not match any known prefix.

For example, Content Delivery Networks (CDN) use anycast [RFC4786] [RFC7094] to improve the quality of service by bringing content closer to users. An anycast IP address is assigned to devices in different locations, and incoming requests are routed to the closest location. Usually, only locations with multiple connectivity announce the anycast IP address through BGP. The CDN server receives requests from users and creates tunnels to the edge locations, where content is sent directly to users using direct server return (DSR). DSR requires servers in the edge locations to use the anycast IP address as the source address in response packets. However, these edge locations do not announce the anycast prefixes through BGP, so an intermediate AS with existing inter-domain SAV mechanisms may improperly block these response packets.

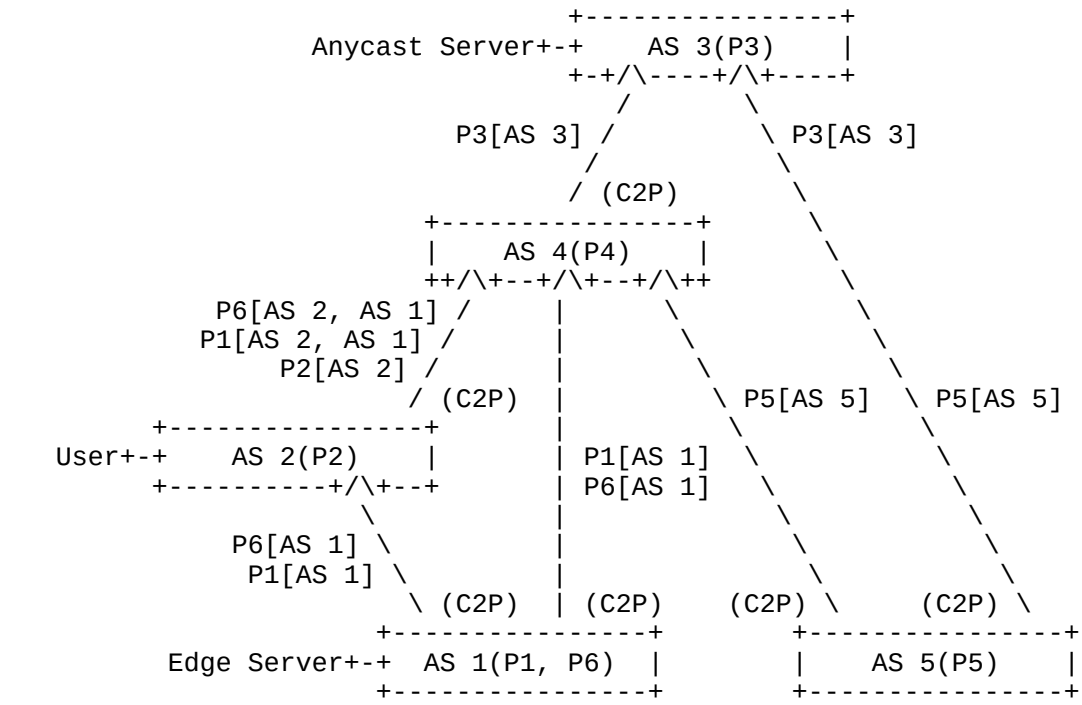

P3 is the anycast prefix and is only advertised by AS 3 through BGP.

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

Figure 3 illustrates a DSR scenario where the anycast IP prefix P3 is only 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 1 and AS 4 have deployed inter-domain SAV, while other ASes have not. When users in AS 2 send requests to the anycast destination IP, the forwarding path is AS

2->AS 4->AS 3. The anycast servers in AS 3 receive the requests and tunnel them to the edge servers in AS 1. Finally, the edge servers send the content to the users with source addresses in prefix P3. The reverse forwarding path is AS 1->AS 4->AS 2. Since AS 4 does not receive routing information for prefix P3 from AS 1, EFP-uRPF with algorithm A/B, and all other existing uRPF-based mechanisms at the customer interface of AS 4 facing AS 1 will improperly block the legitimate response packets from AS 1.

Moreover, EFP-uRPF with algorithm B may also permit spoofing traffic improperly in scenarios where source address spoofing within a customer cone occur. We provide illustrations of these scenarios using an example in the following. The source address spoofing within a customer cone represents a class of scenario where spoofing traffic comes from a customer AS within a customer cone and the spoofed source addresses belong to this customer cone.

4.1.3. Source Address Spoofing within a Customer Cone

Figure 4 portrays a scenario of source address spoofing within a customer cone and is used to analyze the gaps of uRPF-based mechanisms below.

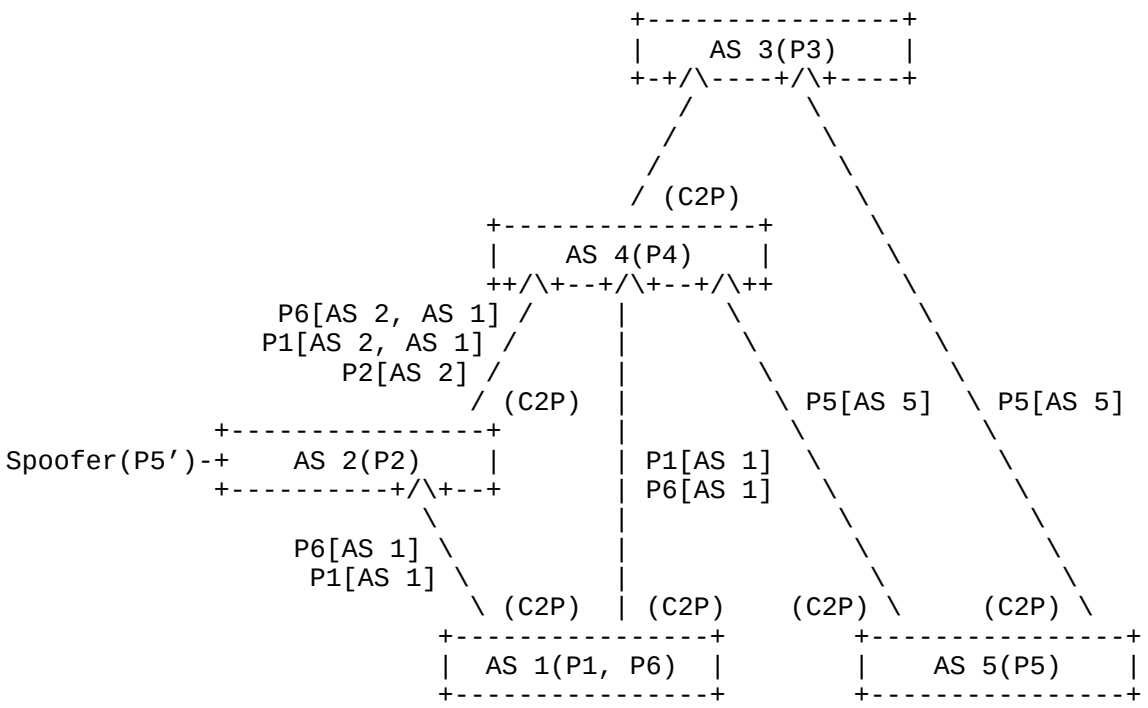

P5' is the spoofed source prefix P5 by the spoofer which is inside of AS 2 or connected to AS 2 through other ASes.

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 with algorithm 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, when the spoofer which is inside of AS 2 or connected to AS 2 through other ASes sends spoofing packets with spoofed source addresses in P5 to AS 3, AS 4 will improperly permit these packets, thus enabling the spoofing traffic to propagate.

In scenarios like these, Strict uRPF, FP-uRPF, VRF uRPF, and EFP-uRPF with algorithm A do not suffer from improper permit problems. This is because these mechanisms enforce strict filtering rules that ensure packets with source addresses in P5 are only permitted to arrive at AS 4's customer interfaces facing AS 5.

4.2. SAV at Provider/Peer Interfaces

SAV is used at provider/peer interfaces to validate traffic entering the customer cone, including both legitimate and spoofing traffic. To prevent packets with spoofed source addresses from the provider/peer AS, ACL-based ingress filtering and/or Loose uRPF can be deployed [nist].

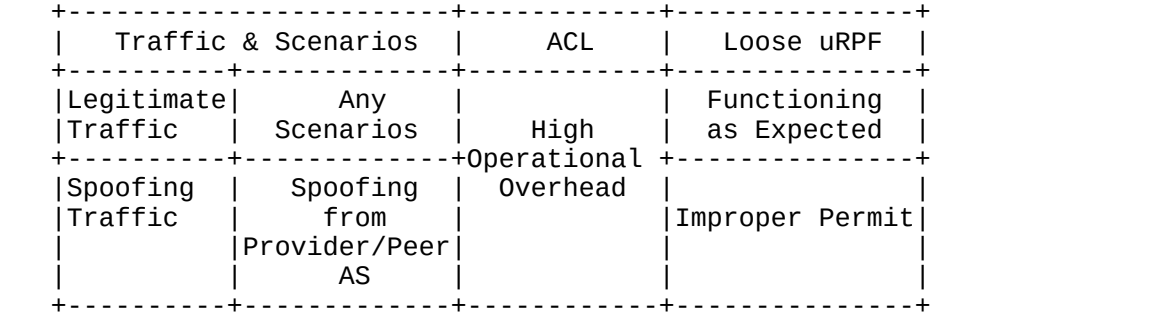

"Spoofing from provider/peer AS" represents a class of scenario where source address spoofing traffic from provider/peer AS occurs and the spoofed source addresses belong to the customer cone which the spoofing traffic enters.
"Functioning as Expected" represents the inter-domain SAV mechanism does not cause improper block for legitimate traffic or improper permit for spoofing traffic in the corresponding scenarios, and has low operational overhead.

Figure 5: The gaps of ACL-based ingress filtering, and Loose uRPF in the corresponding scenarios.

Figure 5 summarizes the gaps of ACL-based ingress filtering and Loose uRPF for SAV at provider/peer interfaces in the corresponding scenarios. ACL-based ingress filtering effectively blocks spoofing traffic from provider/peer AS, while appropriately allowing legitimate traffic. However, these methods may come with 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 the following, we expose the limitations of ACL-based ingress filtering and Loose uRPF for SAV at provider/peer interfaces in scenarios of source address spoofing from provider/peer AS. The source address spoofing from provider/peer AS represents a class of scenario where spoofing traffic comes from a provider/peer AS and the spoofed source addresses belong to the customer cone which the spoofing traffic enters.

4.2.1. Source Address Spoofing from Provider/Peer AS

Figure 6 depicts the scenario of source address spoofing from provider/peer AS and is used to analyze the gaps of ACL-based ingress filtering and Loose uRPF below.

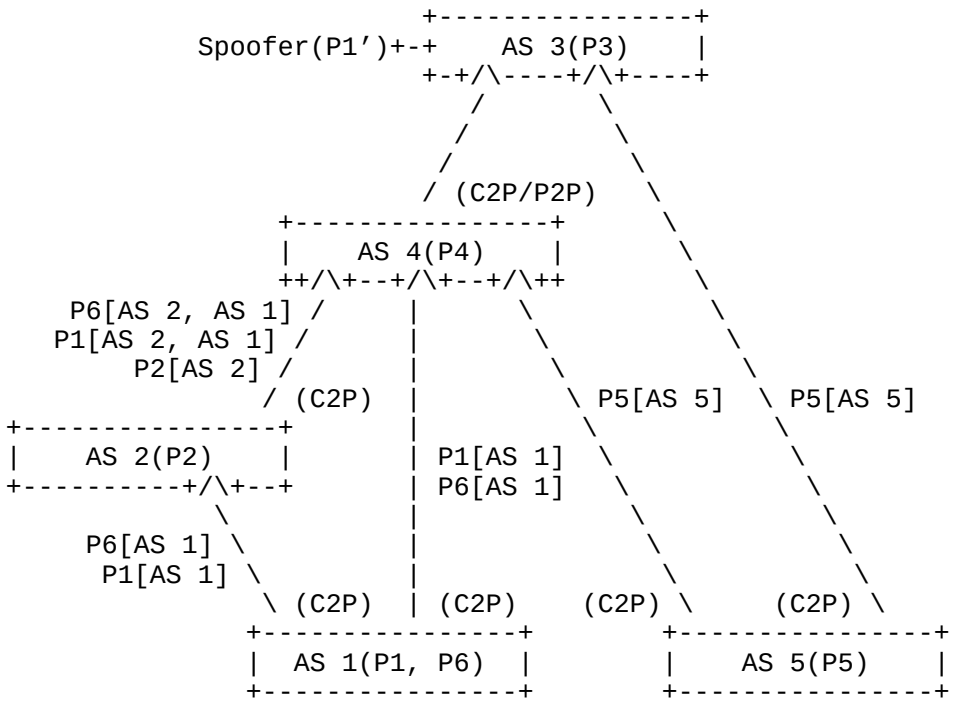

P1' is the spoofed source prefix P1 by the spoofer which is inside of AS 3 or connected to AS 3 through other ASes.

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

In the case of 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. The arrows in Figure 6 represent the commercial relationships between

ASes. AS 3 acts as the provider or lateral peer of AS 4 and the provider for AS 5, while AS 4 serves 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.

By applying ACL-based ingress filtering at the provider/peer interface of AS 4, the ACL rules can block any packets with spoofed source addresses from AS 3 in P1. However, this approach incurs heavy operational overhead, as it requires network operators to update the ACL rules promptly based on changes in prefixes or topology of AS 4’s customer cone. Otherwise, it may cause improper block of legitimate traffic or improper permit of spoofing traffic.

Loose uRPF can greatly reduce the operational overhead because it uses the local FIB as information source, and can adapt to changes in the network. However, it would improperly permit spoofed packets. In Figure 6, Loose uRPF is enabled at AS 4’s provider/peer interface, while EFP-uRPF is enabled at AS 4’s customer interfaces. A spoofer inside AS 3 or connected to it through other ASes may send packets with source addresses spoofing P1 to AS 2. As AS 3 lacks deployment of inter-domain SAV, the spoofing packets will reach AS 4’s provider/peer interface. With Loose uRPF, AS 4 cannot block them at its provider/peer interface facing AS 3, and thus resulting in improper permit.

5. Problem Statement

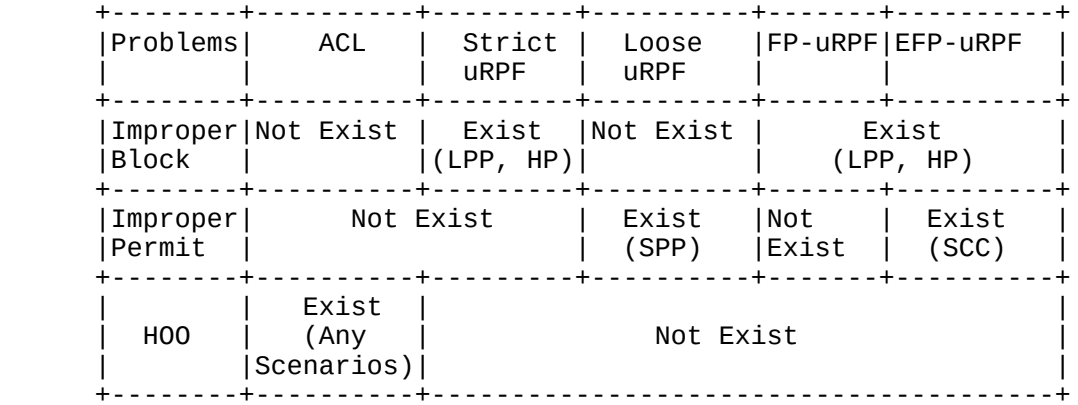

HOO: High Operational Overhead.
"LPP" represents a class of scenario called limited propagation of prefixes.
"HP" represents a class of scenario called hidden prefixes.
"SPP" represents a class of scenario called source address spoofing from provider/peer AS.
"SCC" represents a class of scenario called source address spoofing within a customer cone.

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.

Based on the analysis above, we conclude that existing inter-domain SAV mechanisms exhibit limitations in asymmetric routing scenarios, leading to potential issues of improper block or improper permit. Additionally, these mechanisms can result in high operational overhead, especially when network routing undergoes dynamic changes. Figure 7 provides a comprehensive summary of 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.

For ACL-based ingress filtering, network operators need to manually update ACL rules to adapt to network changes. Otherwise, they may cause improper block or improper permit issues. Manual updates induce high operational overhead, especially in networks with frequent policy and route changes.

Strict uRPF and Loose uRPF are automatic SAV mechanisms, thus they do not need any manual effort to adapt to network changes. However, they have issues in scenarios with asymmetric routing. Strict uRPF may cause improper block problems when an AS is multi-homed and

routes are not symmetrically announced to all its providers. This is because the local FIB may not include the asymmetric routes of the legitimate packets, and Strict uRPF only uses the local FIB to check the source addresses and incoming interfaces of packets. Loose uRPF may cause improper permit problems and fail to prevent source address spoofing. This is because it is oblivious to the incoming interfaces of packets.

FP-uRPF improve Strict uRPF in multi-homing scenarios. However, they still have improper block issues in asymmetric routing scenarios. For example, they may not handle the cases of limited propagation of prefixes. These mechanisms use the local RIB to learn the source prefixes and their valid incoming interfaces. But the RIB may not have all the prefixes with limited propagation and their permissible incoming interfaces.

EFP-uRPF allows the prefixes from the same customer cone at all customer interfaces. This solves the improper block problems of FP-uRPF in multi-homing scenarios. However, this approach also compromises partial protection against spoofing from the customer cone. EFP-uRPF may still have improper block problems when it does not learn legitimate source prefixes. For example, hidden prefixes are not learned by EFP-uRPF.

Finally, existing inter-domain SAV mechanisms cannot work in all directions (i.e. interfaces) of ASes to achieve effective SAV. Network operators need to carefully analyze the network environment and choose appropriate SAV mechanism for each interface. This leads to additional operational and cognitive overhead, which hinders the rate of adoption of inter-domain SAV.

6. Requirements for New Inter-domain SAV Mechanisms

This section lists the requirements which can help bridge the technical gaps of existing inter-domain SAV mechanisms. These requirements serve as the practical guidelines that can be met, in part or in full, by proposing new techniques.

6.1. Accurate Validation

The new inter-domain SAV mechanism MUST improve the validation accuracy in all directions of ASes over existing inter-domain SAV mechanisms, while working in incremental/partial deployment and providing necessary security guarantee.

6.1.1. Improving Validation Accuracy over Existing Mechanisms

The new inter-domain SAV mechanism MUST avoid improper blocking and reject more spoofed traffic than existing inter-domain SAV mechanisms. To achieve this, for an AS performing inter-domain SAV on an interface connected to a neighboring AS, it MUST permit all prefixes whose legitimate traffic (using them as source addresses) can reach that interface, while blocking all other prefixes that cannot. This general principle applies to customer, lateral peer, and provider interfaces. Multiple sources of SAV-related information, such as ROA and ASPA objects, BGP Update data, SAV-specific information, and management information can be leveraged to meet this requirement.

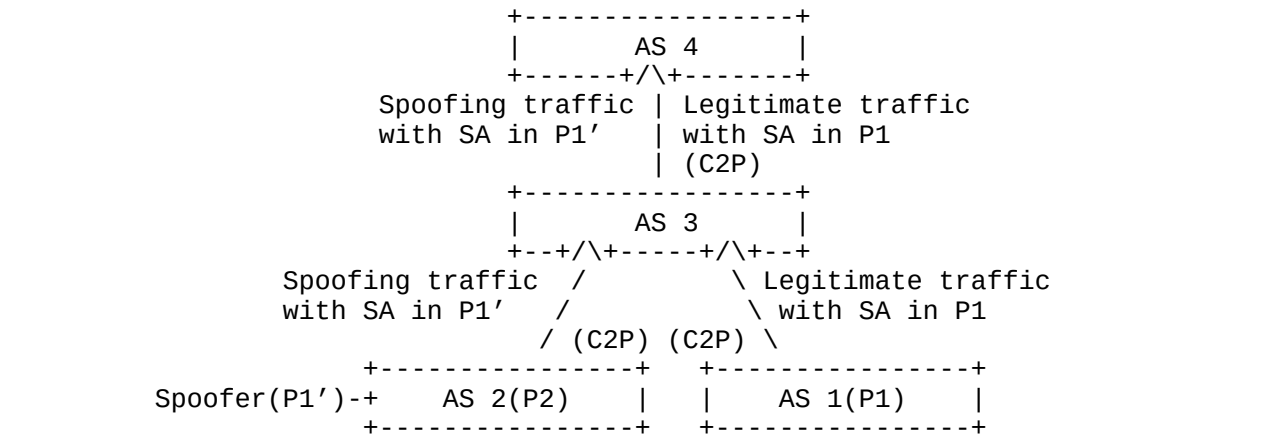

P1' is the spoofed source prefix P1 by the spoofer which is inside of AS 2 or connected to AS 2 through other ASes.
AS 4 performs SAV at its interface facing AS 3.
The legitimate traffic with SA in P1 arrives at AS 4 along the path AS 1->AS 3->AS 4.
The spoofing traffic with SA in P1' arrives at AS 4 along the path AS 2->AS 3->AS 4.

Figure 8: An example where both spoofing and legitimate traffic arrive from the same direction.

The path taken by the traffic with spoofed source address (i.e., spoofed traffic) may overlap with a path for the legitimate traffic. Such scenarios could result in improper permit of the spoofed traffic at the AS doing SAV unless an AS located at or prior to the merging point of the overlap is also performing inter-domain SAV. As illustrated in Figure 8, both spoofed and legitimate traffic traverse the same link between AS 3 and AS 4. In this case, SAV filtering at AS 4's interface facing AS 3 cannot differentiate between the two.

The spoofed traffic in such scenarios is incrementally mitigated (i.e., blocked) with the wider deployment of SAV. For example, AS 3 can deploy SAV on its interfaces facing AS 1 and AS 2 to facilitate blocking of the spoofed traffic while admitting and propagating the legitimate traffic.

6.1.2. Working in Incremental/Partial Deployment

The new inter-domain SAV mechanism MUST NOT assume pervasive adoption (including the adoption of both SAV and SAV-related information) and SHOULD benefit early adopters by providing effective protection from spoofing of source addresses even when it is partially deployed in the Internet. Not all AS border routers can support the new SAV mechanism at once, due to various constraints such as capabilities, versions, or vendors. The new SAV mechanism SHOULD NOT be less effective than existing mechanisms in its capability of protection from source address spoofing for any type of peering interface (customer, lateral peer, and provider) even under partial deployment.

6.1.3. Providing Necessary Security Guarantee

The new inter-domain SAV mechanism SHOULD secure the communicated SAV-specific information between ASes and prevent malicious ASes from generating forged information.

6.2. Automatic Update

The new inter-domain SAV mechanism SHOULD update SAV rules and detect the changes of SAV-specific information automatically while guaranteeing convergence.

6.2.1. Reducing Operational Overhead

The new inter-domain SAV mechanism MUST be able to adapt to dynamic networks and asymmetric routing scenarios automatically, instead of relying on manual update. At least, it MUST have less operational overhead than ACL-based ingress filtering.

6.2.2. Guaranteeing Convergence

The new inter-domain SAV mechanism SHOULD promptly detect the network changes and launch the convergence process quickly. It is essential that the new inter-domain SAV mechanism converges towards accurate SAV rules in a proper manner, effectively reducing improper block and improper permit throughout the whole convergence process.

7. Inter-domain SAV Scope

The new inter-domain SAV mechanisms should work in the same scenarios as existing ones. Generally, it includes all IP-encapsulated scenarios:

* Native IP forwarding: This includes both global routing table forwarding and CE site forwarding of VPN.

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

* Both IPv4 and IPv6 addresses.

Scope does not include:

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

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

8. Security Considerations

SAV rules can be generated based on route information (FIB/RIB) or non-route information. If the information is poisoned by attackers, the SAV rules will be false. Legitimate packets may be dropped improperly or malicious traffic with spoofed source addresses may be permitted improperly. Route security should be considered by routing protocols. Non-route information, such as RPKI ASPA objects, should also be protected by corresponding mechanisms or infrastructure. If SAV mechanisms or protocols require exchanging specific information between ASes, some considerations on the avoidance of message alteration or message injection are needed to propose.

The SAV procedure referred in this document modifies no field of packets. So, security considerations on the data plane are not in the scope of this document.

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

[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>.

[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>.

[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>.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <https://www.rfc-editor.org/rfc/rfc6811>.

[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>.

[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>.

11.2. Informative References

[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>.

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

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

[nist] NIST, "Border Gateway Protocol Security and Resilience", 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>.

[bar-sav] NIST, Akamai, "Source Address Validation Using BGP UPDATEs, ASPA, and ROA (BAR-SAV)", 2024, <https://datatracker.ietf.org/doc/draft-ietf-sidrops-bar-sav/>.

Acknowledgements

Many thanks to Jared Mauch, Barry Greene, Fang Gao, Anthony Somerset, Yuanyuan Zhang, Igor Lubashev, Alvaro Retana, Joel Halpern, Aijun Wang, Michael Richardson, Li Chen, Gert Doering, Mingxing Liu, John O’Brien, Roland Dobbins, etc. for their valuable comments on this document.

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: ksriram@nist.gov