SAVNET
Internet-Draft
Intended status: Informational
Expires: 29 November 2026

D. Li
Tsinghua University
L. Qin
L. Liu
Zhongguancun Laboratory
M. Huang
Huawei
K. Sriram
USA NIST
28 May 2026

Problem Statement, Gap Analysis, and Requirements for Inter-Domain Source Address Validation
draft-ietf-savnet-inter-domain-problem-statement-17

Abstract

This document analyzes the problem space and provides a gap analysis of existing inter-domain source address validation (SAV) mechanisms. Based on these findings, it outlines the technical requirements for future 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 29 November 2026.

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 analyzes the problem space and provides a gap analysis of existing inter-domain source address validation (SAV) mechanisms. Based on these findings, it outlines the technical requirements for future improvements. The corresponding work related to intra-domain SAV is documented in [I-D.ietf-savnet-intra-domain-problem-statement], which includes SAV for hosts and customers (non-AS) connected to AS [SAC-004].

In performing inter-domain SAV, the AS validates the source addresses of data traffic received from a neighboring AS, whether that traffic originated within the neighbor's network or is being transited through it. Inter-domain SAV is applied to incoming traffic on external router interfaces directly connected to a neighboring AS. This includes cases where the neighboring AS uses either a public ASN or a private ASN.

The SAV performing AS and the neighbor AS in consideration are connected using external BGP (eBGP). The eBGP sessions include Customer-to-Provider (C2P), Provider-to-Customer (P2C), lateral peering (P2P), and Route Server (RS) 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 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. This document analyzes the available methods and attempts to answer: (1) what are the technical gaps (Section 4), (2) what are the outstanding problems (Section 4.4), and (3) what are the practical requirements for the solutions to these problems (Section 5).

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

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

To address these problems, this document specifies (Section 5) the following key technical requirements for any new solution:

* Improved SAV accuracy over existing mechanisms: Any new inter-domain SAV mechanism must provide improved SAV accuracy in terms of improper block and improper permit over existing mechanisms. It must seek to achieve zero improper blocking (i.e., avoid false positives) in certain scenarios of interest (Section 4). Further, it must improve the directionality of filtering (i.e., achieve greater rejection of spoofed traffic) over existing mechanisms.

* Reduced operational overhead: Any new inter-domain SAV mechanism should be able to automatically detect changes in the SAV-related information (Section 2) and/or SAV-specific information (Section 2) required for generating the SAV list, obtain the updated information, and use the updated information to generate or update the SAV list.

* Benefit in incremental/partial deployment: Any new inter-domain SAV mechanism must not assume pervasive adoption of the SAV method, the SAV-related information, or the SAV-specific information. It should benefit early adopters by providing effective protection from spoofing of source addresses even in partial deployment.

* Providing necessary security guarantee: If any new inter-domain SAV mechanism introduces or uses SAV-specific information, security mechanisms must exist to prevent malicious injection or alteration of the SAV-specific information.

* Efficient convergence: Any new inter-domain SAV mechanism should achieve efficient convergence of the SAV list after any relevant changes occur in the SAV-related information or SAV-specific information used by the mechanism.

Note that this document focuses on inter-domain SAV mechanisms that validate and filter packets without modifying data plane packets (Section 6). This scope limitation is intentional, since allowing packet modification would introduce additional design, forwarding, interoperability, and deployment considerations beyond the problem space studied in this document. Therefore, SAV mechanisms based on data packet modification are outside the scope of this document.

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. (The terms ’improper block’ and ’false positive’ are used synonymously.)

Improper Permit:
: The validation results in packets with spoofed source addresses being permitted improperly due to an inaccurate SAV list. (The terms ’improper permit’ and ’false negative’ are used synonymously.)

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.

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

SAV-related Information:

Routing information (e.g., RIB and FIB) and objects published in the Resource Public Key Infrastructure (RPKI) that were originally proposed for non-SAV purposes but may also be used for SAV. The RPKI objects include existing RPKI object types (e.g., ROAs and ASPAs) as well as any new types that may be proposed.

SAV-specific Information:
Information dedicated to SAV, which may be defined and exchanged between ASes using potentially new inter-AS communication protocol or an extension of an existing protocol. The information may also take the form of new RPKI object type(s) or management information from operators.

Direct Server Return (DSR):
A traffic delivery model commonly used by Content Delivery Networks (CDNs) that use anycast service addresses while delivering data from edge locations that do not announce those addresses. In such deployments, a request is received by the anycast server or location, but the response is sent directly by another server (i.e., the edge location) with the anycast service address as the source address, rather than the address used to reach the edge server. This can create a legitimate hidden-prefix scenario.

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, when not automated, 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 denylist on a provider interface to block source prefixes that are clearly invalid in the inter-domain routing context, such as internal-use-only prefixes of the SAV-performing AS, IANA special purpose prefixes, and unallocated IPv4/IPv6 prefixes. But it is impractical to store and maintain a very large and dynamically varying set of unallocated IPv6 prefixes. Instead, it may be more practical, for example, to compute an ACL denylist containing the internal-use-only prefixes and prefixes originated exclusively by

the SAV-performing AS and subtract the ACL from an allowlist computed by a uRPF method. Also, for the interfaces with a customer AS, the ACL-only method is impractical while other techniques (using uRPF as described below) are more effective. ACL-based ingress SAV filtering has applicability in scenarios such as (1) directly connected subnets with hosts, or (2) broadband cable, fiber-optic cable, or digital subscriber access loop (DSL) access networks. In these cases, where the service provider should have a clear knowledge of IP address prefixes allocated to manage those services, the ACL-only method in an allowlist form is viable.

* 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 is not 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 did not 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

The inadequacies of inter-domain SAV mechanisms can be characterized along three dimensions: improper block (false positives), improper permit (false negatives), and operational overhead. 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 CC. One example of LPP scenarios 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 scenario occurs by selectively propagating different sets of prefixes to different upstream providers. The Hidden Prefix 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 CC scenario arises when a prefix at one 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 CC of the interface in consideration, improper permit is not fully preventable in the case of source address spoofing within a CC.

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 Limited Propagation of a Prefix, Hidden Prefix, and source address spoofing within a CC 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.

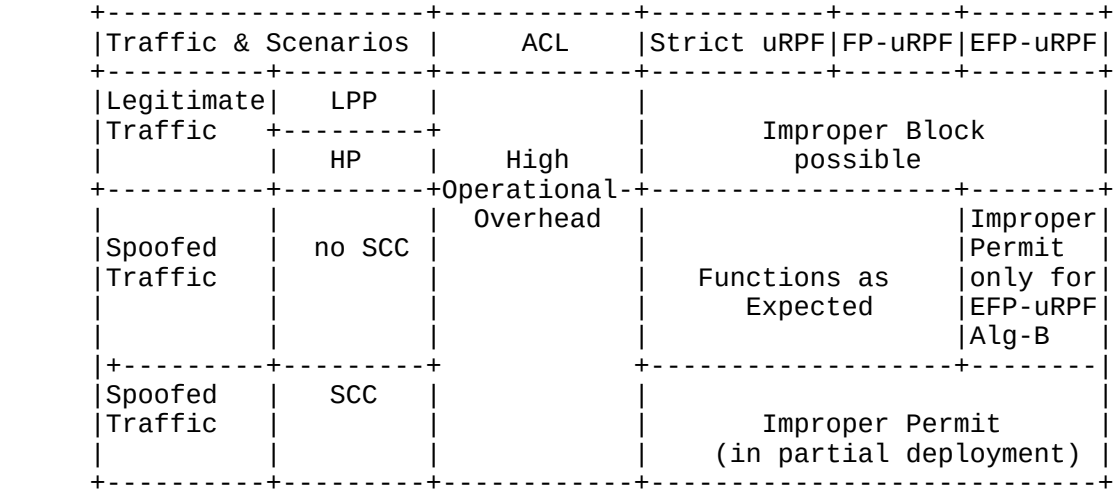

LPP = Limited Propagation of a Prefix
HP = Hidden Prefix
SCC = Spoofing within a CC
'Functions as Expected' connotes the absence of improper permit.
It also connotes low operational overhead.

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 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. 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 of 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.

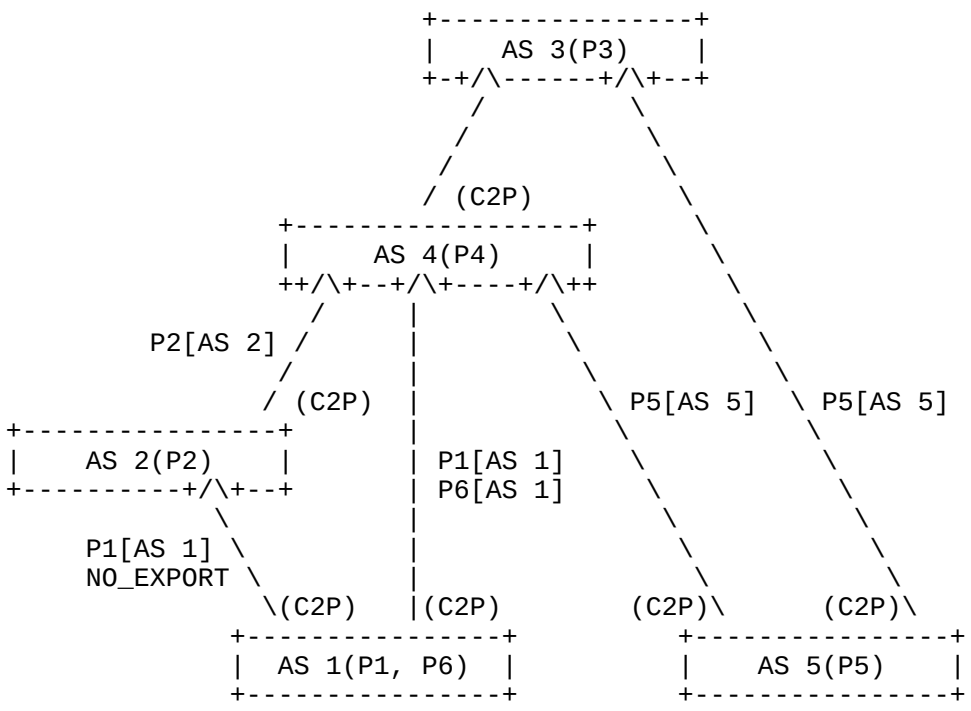

Figure 2: Limited propagation of a prefix caused by NO_EXPORT.

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 prefix 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 Traffic Engineering (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 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 P7 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 P7. The forwarding path for the content packets is AS 1->AS 2. Since AS 2 does not receive routing information for prefix P7 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 2 facing AS 1 will improperly block the response packets from AS 1.

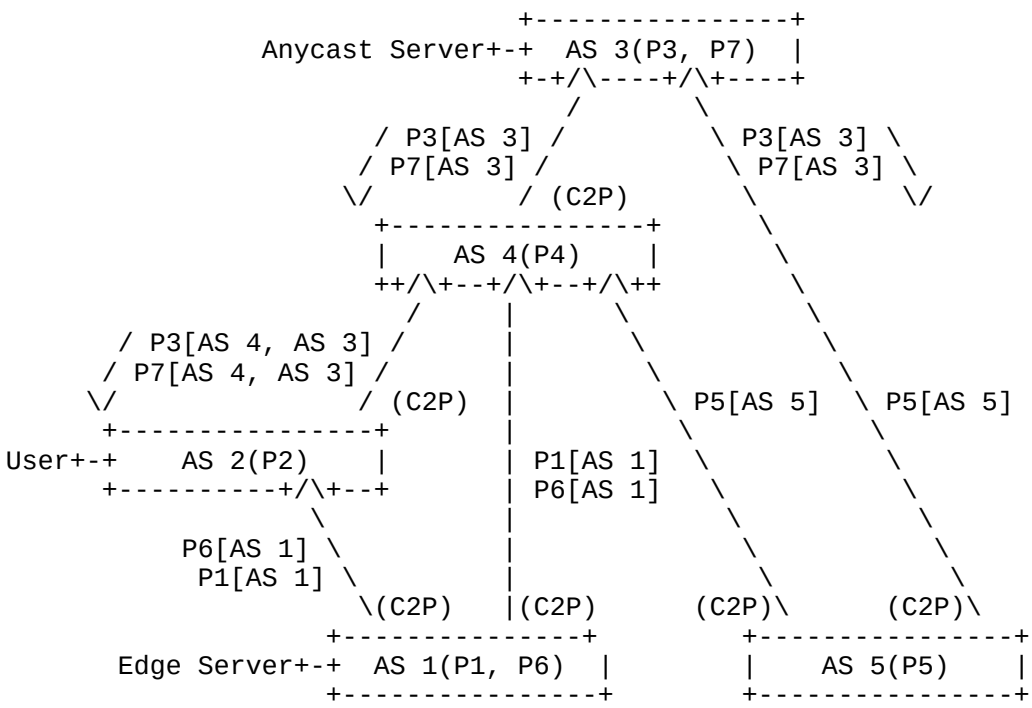

P7 is the anycast prefix and is originated only by AS 3 via BGP. Note that the prefix route propagations relevant to the DSR scenario are depicted; not all prefix propagations are depicted.

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 Scenario

In general, improper permit of spoofed packets in source address spoofing within a CC 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 does not enforce SAV. Now as an example of source address spoofing within a CC, 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 a source address spoofing within a CC 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 not source address spoofing within a CC from the perspective of an individual customer interface of AS 4, but it is source address spoofing within a CC from the perspective of AS 4 looking across all its customer interfaces. EFP-uRPF Alg-B relaxes directionality to reduce (or eliminate) false positives and that makes it more susceptible to source address spoofing within a CC (per the latter perspective). This is expected because EFP-uRPF Alg-B somewhat conservatively applies the same relaxed SAV list across all customer interfaces.

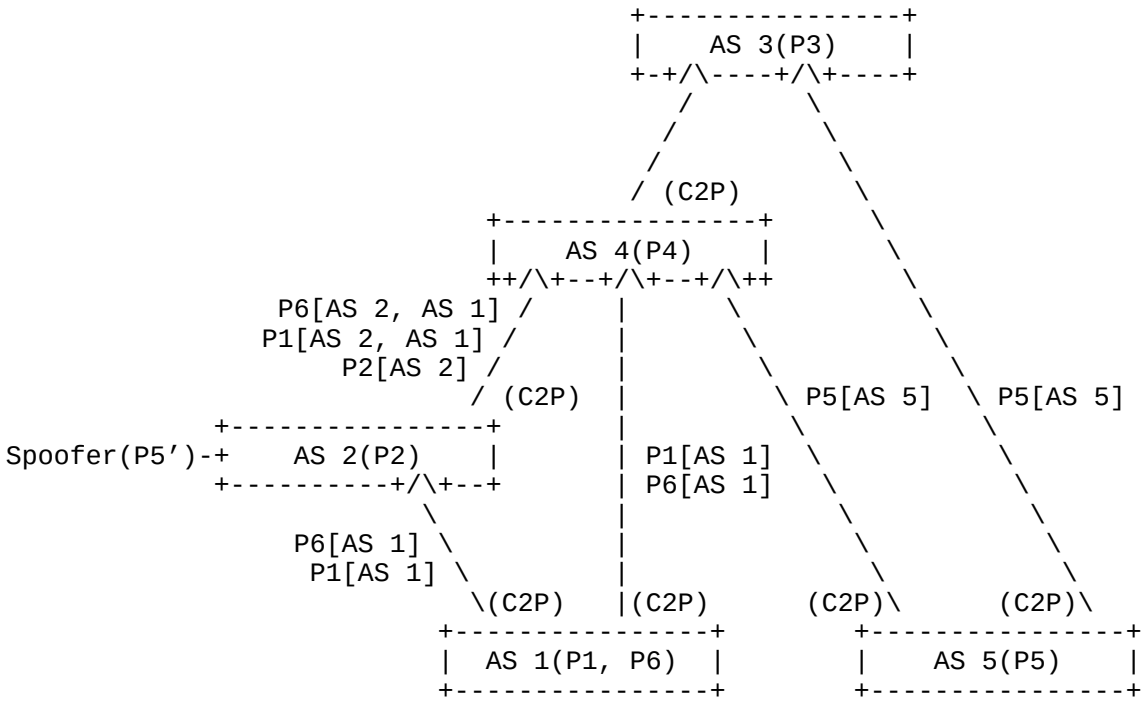

P5' is the spoofed source prefix P5 by the spoofer which is inside of AS 2 or connected to AS 2 through other ASes.

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 CC, 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 a 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 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 customer interfaces may also be used on peer interfaces, 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 a 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, spoofing from providers refers to a scenario in which spoofed traffic comes from provider ASes, either because they originated it or because they forwarded the spoofed traffic that propagated from their neighbor ASes. 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).

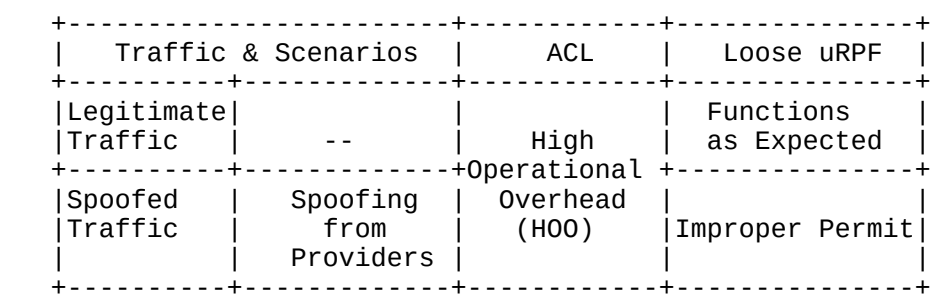

’Functions as Expected’ connotes the absence of improper block. It also connotes low operational overhead.

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 spoofing from providers and is used to analyze the gaps of ACL-based ingress filtering and Loose uRPF.

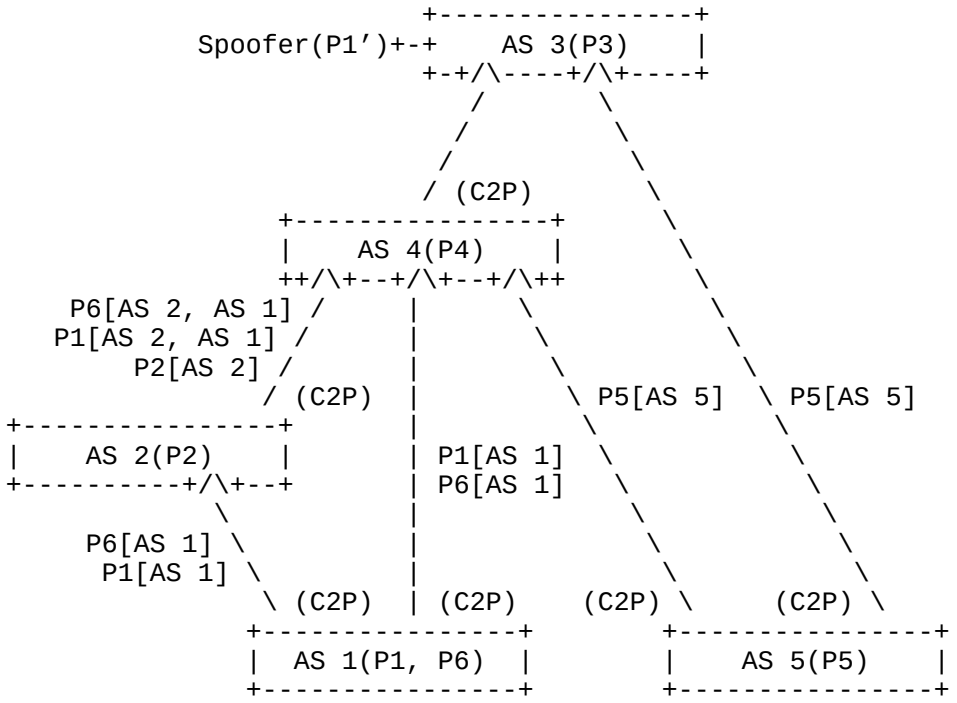

P1' is the spoofed source prefix P1 by the spoofer which is inside of AS 3 or connected to AS 3 through other ASes.

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 an ACL-only method in the form of a denylist at the provider interface of AS 4 (facing AS 3) is impractical (incurs a very high operational overhead) as mentioned in Section 3.

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.

4.4. Gap Analysis Summary

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 Figure 7 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 (Spoofing from Providers)	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
** 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.

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 selective-export 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 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 AS in the CC is spoofed from another AS in the same CC.)

* High operational overhead: ACL-based ingress SAV filtering, when not automated, introduces significant operational overhead, as it needs to update ACL rules manually to adapt to prefix or routing changes in a timely manner. The high operational overhead 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 for SAV-specific information (Section 2), it would be possible to develop comprehensive new SAV algorithms or mechanisms to overcome the existing gaps.

5. 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.

5.1. Accurate Validation

Any new inter-domain SAV mechanism must provide improved SAV accuracy in terms of improper block and improper permit over existing mechanisms. It must seek to achieve zero improper blocking (i.e., avoid false positives) in certain scenarios of interest (Section 4). Further, it must improve the directionality of filtering (i.e., achieve greater rejection of spoofed traffic) over existing mechanisms. The requirement applies for all directions of AS peering (customer, provider, and peer).

5.2. Reducing Operational Overhead

Any new inter-domain SAV mechanism should be able to automatically detect changes in the SAV-related information (Section 2) and/or SAV-specific information (Section 2) required for generating the SAV list, obtain the updated information, and use the updated information to generate the SAV list.

5.3. Early Adopters Benefit in Incremental/Partial Deployment

Any new inter-domain SAV mechanism must not assume pervasive adoption of the SAV method, the SAV-related information, or the SAV-specific information. It should benefit early adopters by providing effective protection from spoofing of source addresses even in partial deployment.

5.4. Providing Necessary Security Guarantee

SAV-related information, e.g., routing information and the existing RPKI signed objects, may be used to design more accurate SAV mechanisms. Such information must be protected during both its creation and dissemination (the BGP security community is already diligent about this). If any new inter-domain SAV mechanism introduces or uses SAV-specific information, security mechanisms must exist to prevent malicious injection or alteration of the SAV-specific information.

5.5. Efficient Convergence

Any new inter-domain SAV mechanism should achieve efficient convergence of the SAV list after any relevant changes occur in the SAV-related information or SAV-specific information used by the mechanism. 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.

6. 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.

SAV mechanisms based on modification of packets in the data plane are outside the scope of this document. Existing architectures or protocols can be inherited by any new SAV mechanisms for greater effectiveness.

7. Security Considerations

The SAV list will be generated based on SAV-related information and/or SAV-specific information. If such information is poisoned by attackers, the resulting SAV list will be inaccurate. Consequently, legitimate traffic may be dropped improperly, or spoofing traffic may be permitted improperly. For SAV mechanisms that use BGP data as input for generating SAV lists, the use of applicable BGP routing security methods is important. Such methods include mechanisms for the prevention, detection, and mitigation of route hijacks, route leaks, and AS_PATH manipulations.

8. IANA Considerations

This document does not request any IANA allocations.

9. Contributors

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

10. References

10.1. Normative References

[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>.

[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>.

[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>.

[I-D.ietf-savnet-intra-domain-problem-statement] Qin, L., Li, D., Wu, J., Huang, M., and N. Geng, "Problem Statement, Gap Analysis, and Requirements for Intra-domain Source Address Validation", Work in Progress, Internet-Draft, draft-ietf-savnet-intra-domain-problem-statement-25, 20 May 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-25>.

10.2. 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>.

[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>.

[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>.

[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>.

[SAC-004] Paul Vixie, "SAC 004 | Security and Stability Advisory Committee - Securing the Edge", 2002, <https://www.icann.org/en/ssac/publications/documents/sac-004-security-and-stability-advisory-committee-securing-the-edge-17-10-2002-en>.

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, Roland Dobbins, Paul Vixie, Amir Herzberg, Jeffrey Haas, and Xueyan Song 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