

SAVNET
Internet-Draft
Intended status: Informational
Expires: 7 April 2026

D. Li
J. Wu
Tsinghua University
L. Qin
M. Huang
Zhongguancun Laboratory
N. Geng
Huawei
4 October 2025

Source Address Validation in Intra-domain Networks Gap Analysis, Problem
Statement, and Requirements
draft-ietf-savnet-intra-domain-problem-statement-19

Abstract

This document provides a gap analysis of existing intra-domain source address validation mechanisms, describes the fundamental problems, and defines the basic 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 7 April 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components

extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	5
1.2. Requirements Language	5
2. Current Operational Intra-domain SAV Mechanisms	5
3. Gap Analysis	6
3.1. Intra-domain SAV for Traffic from Non-BGP Customer Networks or Directly Connected Hosts	6
3.1.1. Asymmetric Routing	7
3.1.2. Hidden Prefix	9
3.2. Intra-domain SAV for Traffic from External ASes	9
4. Problem Statement	10
5. Requirements for New SAV Mechanisms	11
5.1. Accurate Validation	12
5.2. Automatic Update	12
5.3. Working in Incremental Deployment	12
5.4. Fast Convergence	12
5.5. Security	12
6. Security Considerations	13
7. IANA Considerations	13
8. Acknowledgements	13
9. References	13
9.1. Normative References	13
9.2. Informative References	13
Authors' Addresses	14

1. Introduction

Source Address Validation (SAV) defends against source address spoofing. Network operators can enforce SAV at the following levels (see [RFC5210]):

- * Within the access network
- * Within the domain (i.e., the autonomous system)
- * Between domains (i.e., autonomous systems)


Access networks have already deployed SAV mechanisms. These mechanisms typically are deployed on switches and prevent hosts from using the source address of another host on the Internet. Mechanisms include:

- * Source Address Validation Improvement (SAVI) Solution for DHCP [RFC7513]
- * IP Source Guard (IPSG) based on DHCP snooping [IPSG]
- * Cable Source-Verify [cable-verify]

Sadly, access-network SAV mechanisms are not universally deployed. Therefore, intra-domain (i.e., intra-AS) SAV or/and inter-domain (i.e., inter-AS) SAV are required.

This document analyzes intra-domain SAV and focuses on deployment at external interfaces for verifying incoming traffic. SAV at internal interfaces is considered out of scope. Within a domain (i.e., an autonomous system), an external interfaces may connect to a set of hosts, a non-BGP customer network, or an external AS. As illustrated in Figure 1, the goals of intra-domain SAV can be summarized as follows:

- * At external interfaces facing hosts or non-BGP customer networks: Prevent them from injecting packets with source addresses they are not authorized to use into the domain
- * At external interfaces facing external ASes: Prevent those ASes from injecting packets with internal-use-only source addresses into the domain



- SAV at interface 'X' prevents hosts from sending packets with unauthorized source addresses (i.e., addresses outside prefix P2).
- SAV at interface '*' prevents the non-BGP customer network from sending packets with unauthorized source addresses (i.e., addresses outside prefix P1).
- SAV at interface '#' prevents the external AS from injecting packets with internal-use-only source addresses (e.g., prefixes P1 and P2).

Figure 1: Goals of intra-domain SAV

Building on the last goal of intra-domain SAV, inter-domain SAV additionally prevents other ASes from injecting packets with other spoofed source addresses into the domain.

This document provides a gap analysis of the current operational intra-domain SAV mechanisms, identifies key problems to solve, and proposes basic requirements for solutions.

1.1. Terminology

Non-BGP Customer Network: A stub network connected to one or more routers of the AS for Internet connectivity. It only originates traffic and does not participate in BGP routing exchanges with the AS.

SAV Rule: The rule in a router that describes the mapping relationship between a source address (prefix) and the valid incoming interface(s). It is used by a router to make SAV decisions.

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.

SAV-specific Information: The information specialized for SAV rule generation.

1.2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Current Operational Intra-domain SAV Mechanisms

Although BCP 38 [RFC2827] and BCP 84 [RFC3704] specify several ingress filtering methods primarily intended for inter-domain SAV, some of these methods have also been applied to intra-domain SAV in operational practice. This section describes the mechanisms currently used to implement intra-domain SAV.

- * Access Control Lists (ACLs) [RFC2827] are SAV filters that check the source address of each packet against a set of permitted or prohibited prefixes. When applied on a router interface, packets that do not match the ACL entries are blocked. ACLs can be deployed on interfaces facing a non-BGP customer network or a set of hosts, permitting only packets with authorized source addresses. They are also commonly used on interfaces facing an external AS to block packets with unacceptable source addresses, such as internal-use-only prefixes. Since ACLs are typically configured and updated manually, timely updates are essential whenever the set of permitted or prohibited prefixes changes.
- * Strict uRPF [RFC3704] provides an automated SAV filter by validating the source address of each packet against the router's local Forwarding Information Base (FIB). A packet is accepted only if (i) the FIB contains a prefix covering the source address, and (ii) the FIB entry's outgoing interface matches the packet's incoming interface. Otherwise, the packet is discarded. Strict uRPF is commonly used to block spoofed packets originating from a directly connected host or non-BGP customer network.
- * Loose uRPF [RFC3704] also relies on the local FIB for validation, but only checks for the presence of a covering prefix. A packet is accepted if the FIB contains a prefix that covers the source address, regardless of the incoming interface. Loose uRPF is typically used to block spoofed packets that use non-routable or non-global source addresses.

Enhanced Feasible Path uRPF (EFP-uRPF) [RFC8704] is an advanced SAV mechanism specifically designed for inter-domain SAV. It enforces source address validation on router interfaces facing customer ASes by leveraging BGP data received from other ASes. EFP-uRPF is not analyzed in this document, as it is outside the scope of intra-domain SAV.

3. Gap Analysis

This section analyzes the gaps and key challenges of the current operational intra-domain SAV mechanisms.

3.1. Intra-domain SAV for Traffic from Non-BGP Customer Networks or Directly Connected Hosts

To achieve the first goal described in Section 1, an AS operator can deploy ACL rules or strict uRPF on the appropriate routers to enforce intra-domain SAV for traffic originating from non-BGP customer networks or directly connected hosts.


For example, an AS operator can configure an ACL on router interfaces facing a non-BGP customer network or directly connected hosts, specifying the set of prefixes authorized for use as source addresses. The router then blocks any packet whose source address falls outside this set. The main drawback of ACL-based SAV is its high operational overhead. Because ACLs are typically maintained manually, operators must update them promptly to reflect changes in prefixes or topology. Failure to do so may result in outdated ACLs that inadvertently block legitimate traffic.

Strict uRPF automatically generates and updates SAV rules, but it may drop legitimate packets in scenarios such as asymmetric routing or hidden prefixes. The following subsections describe two specific gap scenarios that arise when using strict uRPF for intra-domain SAV.

3.1.1. Asymmetric Routing

Asymmetric routing means a packet traverses from a source to a destination in one path and takes a different path when it returns to the source. Asymmetric routing can occur within an AS due to routing policy, traffic engineering, etc. For example, a non-BGP customer network connected to multiple routers of the AS may need to perform load balancing on incoming traffic, thereby resulting in asymmetric routing.

Figure 2 illustrates an example of asymmetric routing. The non-BGP customer network owns prefix 2001:db8::/56 [RFC6890] and connects to two routers of the AS, Router 1 and Router 2. Router 1, Router 2, and Router 3 exchange routing information via the intra-domain routing protocol. To achieve load balancing for inbound traffic, the non-BGP customer network expects traffic destined for 2001:db8:0::/56 to enter through Router 1, and traffic destined for 2001:db8:0:100::/56 to enter through Router 2. To this end, Router 1 advertises 2001:db8:0::/56 and Router 2 advertises 2001:db8:0:100::/56 through the intra-domain routing protocol. Figure 2 also shows the corresponding FIB entries of Router 1 and Router 2 for the two prefixes.



FIB of Router 1		FIB of Router 2	
Dest	Next_hop	Dest	Next_hop
2001:db8:0::/56	Non-BGP Customer Network	2001:db8:0:100::/56	Non-BGP Customer Network
2001:db8:0:100::/56	Router 3	2001:db8:0::/56	Router 3

The legitimate traffic originated from non-BGP customer network with source addresses in 2001:db8:0:100::/56 will be improperly blocked by strict uRPF on Router 1.

Figure 2: An example of asymmetric routing

While the non-BGP customer network does not expect traffic destined for the prefix 2001:db8:0:100::/56 to arrive via Router 1, it can still send traffic with source addresses within 2001:db8:0:100::/56 to Router 1. As a result, data packets between the non-BGP customer network and Router 1 may follow asymmetric paths. Arrows in the figure indicate the direction of traffic flow.

If Router 1 enforces strict uRPF by checking the FIB entry for the prefix 2001:db8:0:100::/56, the corresponding SAV rule would only allow packets with a source address from 2001:db8:0:100::/56 that arrive via Router 3. Consequently, when the non-BGP customer network sends packets with a source address in 2001:db8:0:100::/56 to Router

1, strict uRPF would incorrectly drop these legitimate packets. Similarly, if Router 2 enforces strict uRPF, it would incorrectly block legitimate packets from the non-BGP customer network that use source addresses within the prefix 2001:db8:0::/56.

3.1.2. Hidden Prefix

The intra-domain hidden prefix scenario refers to two situations in which a host or non-BGP customer legitimately originates traffic using source addresses that are not visible to the intra-domain routing protocol:

- * A host (for example, a cloud server instance operated by a tenant) that originates traffic with a source address not allocated by the AS operator, for legitimate purposes such as Direct Server Return (DSR) deployments.
- * A non-BGP customer network that originates traffic with a source address not advertised to the AS operator, also for valid operational reasons.

For ACL-based SAV, enforcing correct filtering in these scenarios requires authoritative information that explicitly specifies which source addresses the host or non-BGP customer is authorized to use. In practice, such authoritative information is often missing.

Existing uRPF-based mechanisms (strict uRPF or loose uRPF) also fail in hidden prefix scenarios. They will drop packets from hidden prefixes because the source addresses are absent from the router's FIB or are received from unexpected interfaces.

3.2. Intra-domain SAV for Traffic from External ASes

To achieve the second goal described in Section 1, intra-domain SAV is typically deployed on router interfaces facing external ASes to block packets carrying internal-use-only source addresses (see Figure 3). ACL-based SAV is commonly used for this purpose. The AS operator can configure ACL rules containing a set of unacceptable prefixes (for example, internal-use-only prefixes) to block any packet with a source address within these prefixes. However, the operational overhead of maintaining ACL rules can be extremely high, particularly when multiple router interfaces require such configurations, as illustrated in Figure 3.




Figure 3: Intra-domain SAV for traffic from external ASes

In addition, loose uRPF can be used in this context to block packets from external ASes that carry non-global or non-routed source addresses. However, it may allow spoofed packets using internal-use-only source addresses, since internal-use-only prefixes exist in the router’s local FIB.

4. Problem Statement

As discussed above, current operational intra-domain SAV mechanisms have significant limitations with respect to automatic updates and accurate validation.

ACL-based SAV relies entirely on manual maintenance, resulting in high operational overhead in dynamic networks. To ensure the accuracy of ACL-based SAV, AS operators must manually update ACL rules whenever prefixes or topology change; otherwise, packets may be improperly blocked or permitted.

Strict uRPF can automatically update SAV rules, but it may block legitimate traffic in the asymmetric routing or hidden prefix scenarios. As discussed in Section 3.1, strict uRPF may mistakenly consider a valid incoming interface as invalid, resulting in legitimate packets being dropped (i.e., an improper block problem).

Loose uRPF is also an automated SAV mechanism, but its rules are overly permissive. As discussed in Section 3.2, any spoofed packet with a source address present in the FIB may be accepted by loose uRPF (i.e., an improper permit problem).

In summary, even if an AS operator has a comprehensive view and can configure correct ACL rules, manual maintenance imposes high operational overhead and may result in improper blocks due to operator oversight. uRPF cannot guarantee the accuracy of SAV because it relies solely on the router's local FIB to determine SAV rules, which may not correspond to the incoming interfaces of legitimate packets. Consequently, strict uRPF may block legitimate traffic in asymmetric routing and hidden prefix scenarios, while loose uRPF has limited effectiveness against source address spoofing, as it only blocks non-global or non-routed addresses. For hidden prefix scenarios, the key challenge remains how to provide authoritative information that allows the host or non-BGP customer to legitimately use such source addresses.

Another consideration is that uRPF-based mechanisms rely on routing information to make SAV decisions, assuming that the routing information in the local FIB is correct. If the routing information is incorrect, SAV decisions may also be incorrect, potentially resulting in improper blocks or permits. It should be emphasized that ensuring the correctness of routing information is the responsibility of mechanisms or operational processes outside the scope of SAV. Network operators and SAV mechanisms are encouraged to leverage such solutions to validate the routing information used by SAV.

5. Requirements for New SAV Mechanisms

This section outlines five general requirements for technical improvements that should be considered when designing future intra-domain SAV architectures and solutions. These informational requirements can not be used to initiate standards-track protocol changes.

5.1. Accurate Validation

The new intra-domain SAV mechanism MUST improve the accuracy of existing intra-domain SAV mechanisms. It MUST achieve the goals described in Section 1, preventing spoofed traffic from entering the domain. At the same time, it MUST avoid blocking legitimate packets, particularly in the presence of prefix changes, asymmetric routes, or hidden prefixes. To overcome the improper block problems, routers may need to use additional information (e.g., SAV-specific information) beyond the local FIB information to make SAV decisions. By integrating such information, routers can account for asymmetric routes and hidden prefixes, resulting in more accurate SAV rules.

5.2. Automatic Update

The new intra-domain SAV mechanism MUST be capable of automatically generating and updating SAV rules on routers, rather than relying entirely on manual updates as in ACL-based SAV. Although some initial configuration may be necessary to improve SAV accuracy, automation reduces the subsequent operational overhead for the AS operator.

5.3. Working in Incremental Deployment

The new mechanism MUST support incremental deployment and MUST provide incremental benefits under such partial deployment. In an incremental deployment scenario, the mechanism MUST avoid improper blocks and MUST clearly specify the extent to which the goals described in Section 1 can be partially achieved.

5.4. Fast Convergence

The new intra-domain SAV mechanism MUST be able to update SAV rules promptly when prefixes, routes, or topology change within an AS. If SAV-specific information is communicated via a protocol, two considerations are essential. First, the mechanism MUST allow routers to learn updated SAV-specific information in a timely manner. Second, the mechanism MUST NOT transmit excessive SAV-specific information, as this could significantly increase the burden on the routers' control planes and potentially degrade the performance of existing protocols.

5.5. Security

The new intra-domain SAV mechanisms MUST NOT introduce additional security vulnerabilities or create confusion in existing intra-domain architectures or protocols. Section 6 details the security scope and considerations for the new intra-domain SAV mechanism.

6. Security Considerations

Similar to the security scope of intra-domain routing protocols, intra-domain SAV mechanisms can ensure the integrity and authentication of protocol messages that convey the required SAV-specific information and can help prevent unintentional misconfigurations. It is not necessary for SAV mechanisms to protect against compromised or malicious intra-domain routers that attempt to poison existing control or management plane protocols. Such compromised or malicious routers could not only affect SAV, but also disrupt the entire intra-domain routing domain. Security mechanisms to defend against these attacks are beyond the scope and capability of intra-domain SAV.

7. IANA Considerations

This document does not request any IANA allocations.

8. Acknowledgements

Many thanks to the valuable comments from: Jared Mauch, Joel Halpern, Aijun Wang, Michael Richardson, Gert Doering, Libin Liu, Li Chen, Tony Przygienda, Yingzhen Qu, James Guichard, Linda Dunbar, Robert Sparks, Stephen Farrel, Ron Bonica, etc.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

9.2. Informative 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/info/rfc2827>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.

- [RFC5210] Wu, J., Bi, J., Li, X., Ren, G., Xu, K., and M. Williams, "A Source Address Validation Architecture (SAVA) Testbed and Deployment Experience", RFC 5210, DOI 10.17487/RFC5210, June 2008, <<https://www.rfc-editor.org/info/rfc5210>>.
- [RFC8704] Sriram, K., Montgomery, D., and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding", BCP 84, RFC 8704, DOI 10.17487/RFC8704, February 2020, <<https://www.rfc-editor.org/info/rfc8704>>.
- [RFC7513] Bi, J., Wu, J., Yao, G., and F. Baker, "Source Address Validation Improvement (SAVI) Solution for DHCP", RFC 7513, DOI 10.17487/RFC7513, May 2015, <<https://www.rfc-editor.org/info/rfc7513>>.
- [cable-verify] "Cable Source-Verify and IP Address Security", January 2021, <<https://www.cisco.com/c/en/us/support/docs/broadband-cable/cable-security/20691-source-verify.html>>.
- [IPSG] "Configuring DHCP Features and IP Source Guard", January 2016, <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_53_se/configuration/guide/2960scg/swdhcp82.html>.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <<https://www.rfc-editor.org/info/rfc6890>>.

Authors' Addresses

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

Jianping Wu
Tsinghua University
Beijing
China
Email: jianping@cernet.edu.cn

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

Mingqing Huang
Zhongguancun Laboratory
Beijing
China
Email: huangmq@mail.zgclab.edu.cn

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