

SAVNET
Internet-Draft
Intended status: Informational
Expires: 17 October 2026

D. Li
J. Wu
Tsinghua University
L. Qin
M. Huang
Zhongguancun Laboratory
N. Geng
Huawei
15 April 2026

Source Address Validation in Intra-domain Networks Gap Analysis, Problem
Statement, and Requirements
draft-ietf-savnet-intra-domain-problem-statement-23

Abstract

This document provides a gap analysis of the current operational intra-domain SAV mechanisms and identifies requirements for new intra-domain SAV solutions.

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 17 October 2026.

Copyright Notice

Copyright (c) 2026 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	4
1.2. Requirements Language	5
2. Problem Statement of Current Operational Intra-domain SAV Mechanisms	5
3. Gap Analysis	6
3.1. Asymmetric Routing Scenario	7
3.2. Hidden Prefix Scenario	9
4. Requirements for New SAV Mechanisms	9
4.1. Accurate Validation	10
4.2. Automatic Updates	11
4.3. Incremental Deployment Support	11
4.4. Fast Convergence	11
4.5. Authentication of Information Used for SAV	11
4.6. Vulnerability Prevention	11
5. Security Considerations	12
6. IANA Considerations	12
7. Acknowledgements	12
8. References	12
8.1. Normative References	12
8.2. Informative References	13
Authors' Addresses	13

1. Introduction

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

- * IP source address validation in the access network
- * IP source address validation at intra-AS/ingress point
- * IP source address validation in the inter-AS Case (neighboring AS)

Some access networks have already deployed SAV mechanisms. These mechanisms typically are deployed on switches in the access network and prevent hosts from using the source address of another host on the Internet [RFC5210]. Mechanisms include:

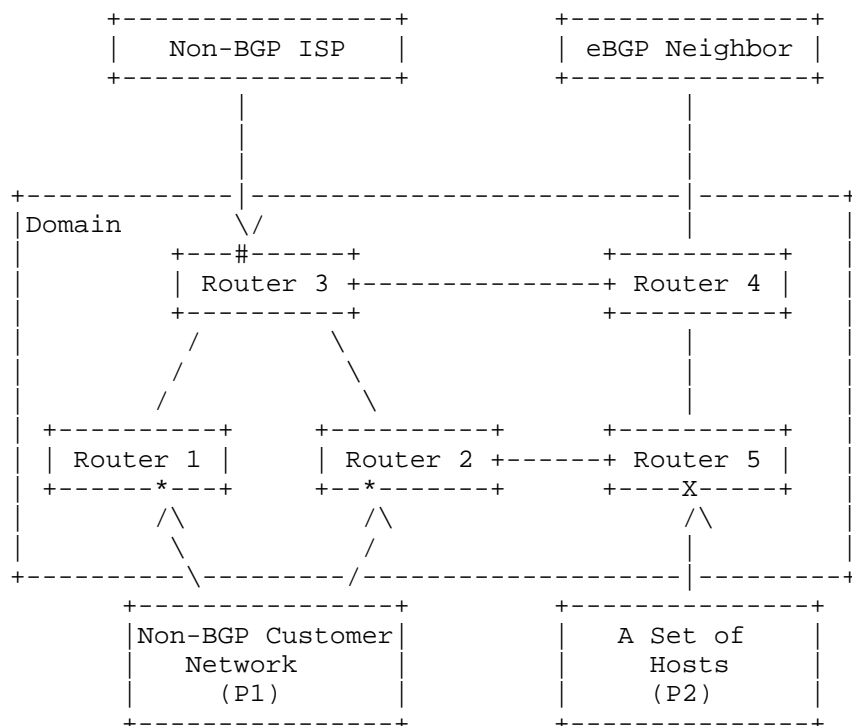
- * Source Address Validation Improvement (SAVI) Solution for DHCP [RFC7513]

- * IP Source Guard (IPSG) based on DHCP snooping [IPSG]
- * Cable Source-Verify [cable-verify]

However, access-network SAV mechanisms are not universally deployed [CAIDA-spoofers]. Therefore, intra-domain (i.e., intra-AS) SAV and inter-domain (i.e., inter-AS) SAV are required [RFC5210].

This document provides a gap analysis of the current operational intra-domain SAV mechanisms and identifies requirements for new intra-domain SAV solutions.

In this document, a domain refers to a routing domain under a single administrative control (e.g., an AS). Intra-domain SAV refers to SAV at a domain's external interfaces that do not carry external BGP (eBGP) sessions (i.e., non-BGP external interfaces). SAV at internal interfaces or BGP-facing external interfaces is considered out of scope. For a domain, as illustrated in Figure 1, a non-BGP external interface may connect to a set of hosts, a non-BGP customer network, or a non-BGP Internet Service Provider (ISP) network. The goal of intra-domain SAV at such interfaces is to prevent traffic using unauthorized source addresses from entering the domain.



This document focuses on SAV at a domain's non-BGP external interfaces including Interfaces 'X', '*', and '#'.

Figure 1: Deployment locations of intra-domain SAV

1.1. Terminology

Non-BGP Customer Network: A stub network (i.e., a network that only originates traffic) connected to its provider network for Internet connectivity and does not participate in eBGP peering with its provider network.

Non-BGP Internet Service Provider (ISP) Network: A network that forwards traffic from its customer network to the Internet and does not participate in eBGP peering with its customer network.

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.

Proper Block: The validation results that packets with spoofed source addresses are blocked by SAV rules.

Proper Permit: The validation results that packets with legitimate source addresses are permitted by 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.

The requirements language is used in Section 4 and applies to implementations of SAV conformant to the listed requirements.

2. Problem Statement of 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 summarizes the problems of mechanisms currently used to implement intra-domain SAV. These mechanisms have significant limitations in terms of automated updates or accurate validation.

- * Access Control Lists (ACLs) can be used as SAV filters [RFC2827] to check the source address of each packet against a set of permitted or prohibited prefixes. When applied on a router interface, each ACL entry (ACE) specifies both matching conditions (e.g., prefixes) and the corresponding action (e.g., permit or deny), and packets are processed accordingly. To ensure correct filtering behavior, changes in SAV state need to be reflected in corresponding ACL rules, which in turn need to be updated in accordance with changes in prefixes or topology; otherwise, packets may be improperly permitted or blocked. In ACL-based

ingress filtering [RFC2827] deployments, maintaining consistency between SAV state and ACL rules can introduce operational challenges, as this update process is often performed manually or requires significant operational intervention.

- * 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. It may block legitimate traffic in the asymmetric routing or hidden prefix scenarios (see Section 3.1 and Section 3.2). Strict uRPF may mistakenly consider a valid incoming interface as invalid, resulting in legitimate packets being blocked (i.e., an improper block problem).
- * 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. Since its rules are overly permissive, any spoofed packet with a source address present in the FIB may be permitted by loose uRPF (i.e., an improper permit problem).
- * Enhanced Feasible Path uRPF (EFP-uRPF) [RFC8704] is an advanced SAV mechanism specifically designed for inter-domain SAV. It enforces SAV on eBGP interfaces facing a customer AS by leveraging BGP data received from external 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.

ACL-based SAV can be deployed on interfaces facing a non-BGP customer network or a set of hosts, permitting only packets with authorized source addresses. Such mechanism can also be applied on interfaces facing a non-BGP ISP network to block packets with prohibited source addresses, including internal-use-only addresses, unallocated addresses, and addresses single-homed to the local domain (e.g., P1 and P2 in Figure 1). A key limitation of ACL-based SAV is the need to maintain consistency between SAV state and ACL rules. Operators need to update ACL rules to reflect changes in prefixes or topology, and delays or inconsistencies in this process may result in outdated rules that inadvertently block legitimate traffic or permit spoofed traffic.

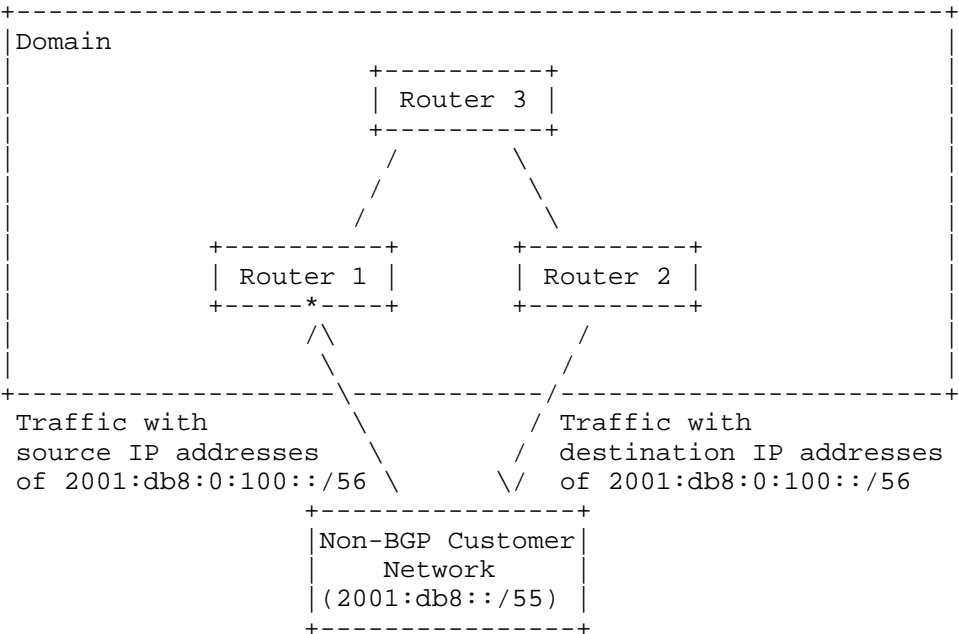
As noted in Section 2.4 of [RFC3704], loose uRPF sacrifices directionality, so its effectiveness in mitigating source address spoofing is very limited, and improper permit problems may occur.

With strict uRPF, 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. Asymmetric Routing Scenario

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::/55 [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

Although the non-BGP customer network does not expect to receive inbound traffic for 2001:db8:0:100::/56 via Router 1, it can send outbound traffic with source addresses in that prefix through 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.

3.2. Hidden Prefix Scenario

The intra-domain hidden prefix scenario refers to 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 within the domain.

- * A host (for example, a cloud server instance operated by a tenant) may originate traffic using a source address not allocated by the AS operator. This can occur in deployments such as Direct Server Return (DSR), where return traffic is sent directly from the server using a service IP address that is not part of the operator's internal routing view.
- * A non-BGP customer network may originate traffic using source addresses that are not advertised to the domain operator. This can occur in scenarios such as Direct Server Return (DSR) deployments or when the customer network uses address space assigned by another provider (e.g., in multi-homing or hybrid connectivity scenarios), and such prefixes are not propagated within the operator's intra-domain routing system.

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.

4. Requirements for New SAV Mechanisms

The limitations described above primarily stem from the lack of SAV-specific authoritative information that can be consistently and automatically consumed by SAV mechanisms. Existing automated uRPF-based approaches derive SAV decisions from routing or forwarding state, which is intended to express reachability rather than authorization of source address usage. As a result, these mechanisms may not provide reliable validation in scenarios such as asymmetric routing or hidden prefixes. In contrast, ACL-based approaches can express source address authorization more precisely, but rely on

ongoing operational intervention, which limits their applicability in dynamic operational environments.

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 blocking or permitting. Ensuring the correctness of routing information is the responsibility of mechanisms or operational processes outside the scope of SAV. However, when SAV relies on routing information or other contextual information, such information is expected to be derived from trusted sources before being used.

This section identifies five requirements to guide the design, development, and evaluation of new intra-domain SAV mechanisms, and to provide a common basis for improving upon existing approaches. These requirements are informational in nature. To avoid misinterpretation or misuse as a normative reference, it is noted that these informational requirements cannot be used to initiate standards-track protocol changes.

4.1. Accurate Validation

Any new intra-domain SAV mechanism **MUST** improve the accuracy of source address validation compared to existing uRPF-based mechanisms. In particular, it **MUST** reduce the occurrence of improper blocks (i.e., blocking legitimate traffic), improper permits (i.e., allowing spoofed traffic), or both. Specifically, it **MUST** satisfy the following conditions:

- * result in fewer improper blocks than strict uRPF, particularly in scenarios involving asymmetric routes or hidden prefixes;
- * result in fewer improper permits than loose uRPF.

To achieve higher SAV accuracy, additional information beyond the local FIB (e.g., SAV-specific information) may be needed to make validation decisions. By integrating such information, routers may have the ability to account for asymmetric routes and hidden prefixes, resulting in more accurate SAV rules.

4.2. Automatic Updates

Any new intra-domain SAV mechanism MUST be capable of automatically collecting and processing relevant information, and using it to derive and update SAV state and corresponding filtering rules on routers. Automation helps reduce operational complexity and maintenance overhead, while allowing some initial configuration to improve SAV accuracy. This ensures the mechanism is deployable in practical networks without introducing excessive management burden.

4.3. Incremental Deployment Support

Any new intra-domain SAV mechanism MUST support incremental deployment and provide measurable benefits even when only a subset of external non-BGP interfaces deploy the mechanism.

4.4. Fast Convergence

If any new intra-domain SAV mechanism requires disseminating SAV-specific information among intra-domain routers via a protocol, two considerations are essential. First, such mechanism MUST allow routers to learn updated SAV-specific information in a timely manner. Second, such mechanism MUST NOT transmit excessive SAV-specific information via a protocol, as this could significantly increase the burden on the routers' control planes and potentially degrade the performance of existing protocols.

4.5. Authentication of Information Used for SAV

Any new intra-domain SAV mechanism MUST use information that is authenticated or trusted, either through verification of its integrity and authenticity, or via an established trust relationship with the information source. If a SAV mechanism introduces new SAV-specific information, such information MUST be authenticated to ensure its integrity and authenticity before being used for SAV decision making.

4.6. Vulnerability Prevention

Any new intra-domain SAV mechanism MUST NOT introduce additional security vulnerabilities to existing intra-domain architectures or protocols. Protection against compromised or malicious intra-domain routers is out of scope, as such routers can compromise not only SAV mechanisms but also the entire intra-domain routing domain.

5. Security Considerations

This document discusses the limitations of existing intra-domain SAV practices and identifies problems and informational requirements for improved intra-domain SAV mechanisms. It does not specify new protocols or mechanisms and, as such, does not introduce any new security considerations.

6. IANA Considerations

This document does not request any IANA allocations.

7. 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, Xueyan Song, etc. We also thank the IETF Directorates and the IESG for their reviews and comments, which helped improve the clarity of this document.

8. References

8.1. 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/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>>.
- [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>>.

8.2. Informative References

- [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>.
- [CAIDA-spoofers] "State of IP Spoofing", n.d., <<https://spoofers.caida.org/summary.php?>>>.
- [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