

SAVNET
Internet-Draft
Intended status: Informational
Expires: 4 December 2026

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

Problem Statement, Gap Analysis, and Requirements for Intra-domain
Source Address Validation
draft-ietf-savnet-intra-domain-problem-statement-26

Abstract

Source address validation (SAV) is an important means to mitigate IP source address spoofing [RFC2827]. This document analyzes the gaps in current operational mechanisms for intra-domain SAV. It also identifies the properties that new intra-domain SAV mechanisms are expected to provide.

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 4 December 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
2. Problem Statement	5
3. Current Operational Intra-domain SAV Mechanisms	6
4. Gap Analysis	6
4.1. Asymmetric Routing Scenario	7
4.2. Hidden Prefix Scenario	9
5. Requirements for New SAV Mechanisms	9
5.1. Accurate Validation	10
5.2. Automatic Updates	10
5.3. Incremental Deployment Support	10
5.4. No Adverse Impact on Routing Convergence and Fast Reroute	11
5.5. Authentication of Information Used for SAV	11
6. Security Considerations	11
7. IANA Considerations	11
8. Acknowledgements	11
9. References	11
9.1. Normative References	11
9.2. Informative References	12
Authors' Addresses	12

1. Introduction

Source Address Validation (SAV) defends against IP source address spoofing [RFC2827]. 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)

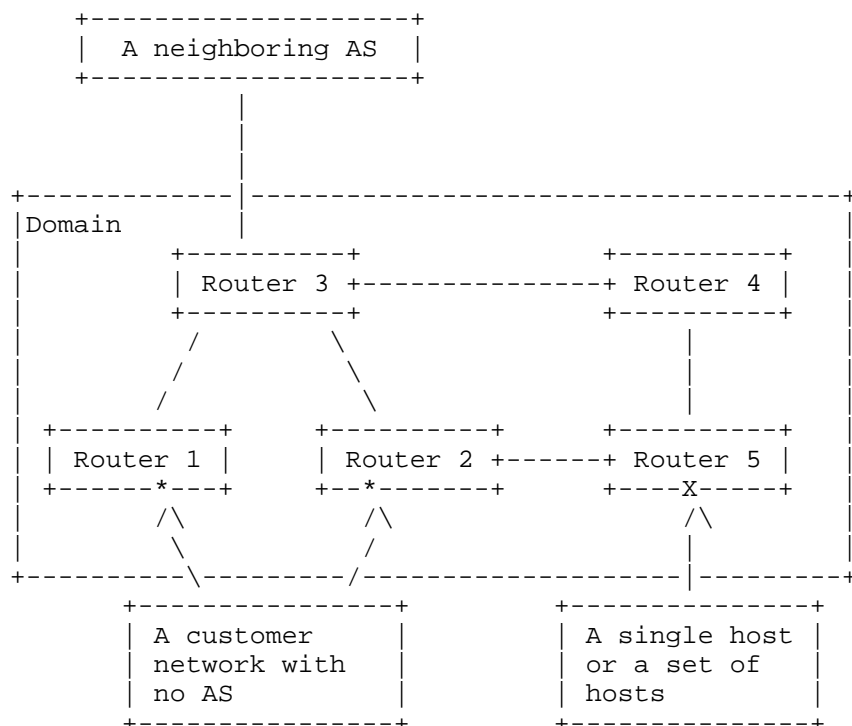
Some access networks have already deployed SAV mechanisms. These mechanisms are typically deployed on switches in the access network and prevent hosts connected to those switches from using the source address of another host on the Internet [RFC5210]. Mechanisms include:

- * 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]. For the purpose of this document, intra-domain SAV is defined as follows:

- * The AS validates the source addresses of data traffic that it originates directly or indirectly. Intra-domain SAV is applied at external interfaces (on routers) facing entities that are not deployed as neighboring ASes and are therefore not covered by inter-domain SAV. For example, as illustrated in Figure 1, an entity can be a single host, a set of hosts, or a customer network with no AS that manages one or more IP prefixes. The entity may source traffic using prefixes assigned by the AS or its own BYOIP prefixes. From the perspective of other ASes, such traffic is originated by the AS.

SAV on traffic received on external interfaces facing a neighboring AS is considered inter-domain SAV, regardless of whether the neighboring AS uses a public ASN or a private ASN. SAV on internal interfaces (e.g., interfaces between Router 1 and Router 3 in Figure 1) is also outside the scope of this document. This is because routers inside the same AS are generally assumed to be trusted, so SAV on internal interfaces provides limited additional benefit when SAV is already applied at external interfaces. In addition, techniques such as fast reroute can make SAV at internal interfaces technically challenging.



Intra-domain SAV is applied at interfaces '*' and 'X'.

Figure 1: Deployment locations of intra-domain SAV

This document analyzes the gaps in current operational mechanisms for intra-domain SAV. It also identifies the properties that new intra-domain SAV mechanisms are expected to provide.

1.1. Terminology

SAV Rule:

The rule that indicates the validity of a specific source IP address or source IP prefix per router interface. It is used by a router to make SAV decisions.

Improper Block:

The validation results in packets with legitimate source addresses being blocked improperly due to inaccurate SAV rules. (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 inaccurate SAV rules. (The terms 'improper permit' and 'false negative' are used synonymously.)

Proper Block:

The validation results in packets with spoofed source addresses being blocked.

Proper Permit:

The validation results in packets with legitimate source addresses being permitted.

SAV-specific Information:

The information specialized for SAV rule generation.

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.

2. Problem Statement

The problems of existing intra-domain SAV mechanisms can be characterized along three dimensions: improper block, improper permit, and operational overhead:

- * Improper block. Existing intra-domain SAV mechanisms may block data packets using legitimate source addresses when the applied SAV rules are inaccurate.
- * Improper permit. Existing intra-domain SAV mechanisms may permit data packets using spoofed source addresses when the applied SAV rules are inaccurate.
- * Operational overhead. Existing intra-domain SAV mechanisms may require operator involvement to determine and update SAV rules. This overhead depends on how much manual effort is needed to keep the SAV rules up to date.

In this document, these three dimensions are used to analyze the gaps in existing intra-domain SAV mechanisms.

3. 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 introduces the mechanisms currently used to implement intra-domain SAV.

- * 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 Access Control Entry (ACE) used for SAV filtering specifies both matching conditions (i.e., prefixes) and the corresponding action (e.g., permit or deny), and packets are processed accordingly.
- * 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.
- * 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.

4. Gap Analysis

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

ACLs can be used on interfaces facing a customer network with no AS or a set of hosts to permit only packets whose source addresses belong to specific prefixes. To ensure correct filtering behavior, the ACLs used for SAV filtering need to be updated when the permitted prefixes change; otherwise, packets may be improperly permitted or blocked. In ACL-based SAV deployments, keeping these ACLs up to date can introduce operational challenges when operators need to detect prefix changes and determine and apply the corresponding ACL updates.

As described in Section 3 and also noted in [RFC3704], loose uRPF sacrifices directionality when validating source addresses of data packets. 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).

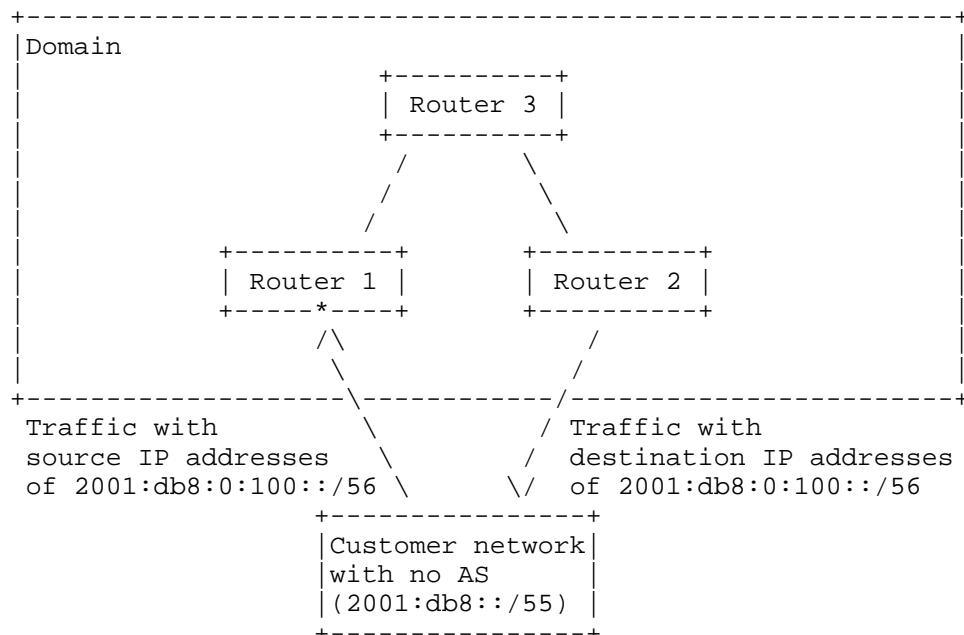
Strict uRPF may block legitimate traffic in the asymmetric routing or hidden prefix scenarios (see Section 4.1 and Section 4.2). It may mistakenly consider a valid incoming interface as invalid, resulting in legitimate packets being blocked (i.e., an improper block problem).

The following subsections describe two specific gap scenarios for intra-domain SAV.

4.1. 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 due to routing policy, traffic engineering, etc.

For example, a customer network with no AS 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 customer network owns prefix 2001:db8::/55 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 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	Customer Network	2001:db8:0:100::/56	Customer Network
2001:db8:0:100::/56	Router 3	2001:db8:0::/56	Router 3

The legitimate traffic originated from the 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 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 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 customer network sends packets with a source address in 2001:db8:0:100::/56 to Router 1, strict uRPF would improperly drop these legitimate packets.

Similarly, if Router 2 enforces strict uRPF, it would improperly block legitimate packets from the customer network that use source addresses within the prefix 2001:db8:0::/56.

4.2. Hidden Prefix Scenario

The intra-domain hidden prefix scenario refers to situations in which a host or a customer network with no AS 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 customer network with no AS may originate traffic using source addresses that are not advertised to the AS 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 the customer network is authorized to use. In practice, such authoritative information is often missing. Strict uRPF and 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.

5. Requirements for New SAV Mechanisms

This section identifies five requirements that can inform the design of new intra-domain SAV mechanisms. These requirements describe the properties that new mechanisms are expected to provide in order to improve upon existing mechanisms, but do not make assumptions about how those properties are achieved. They do not mandate or justify any specific extension to routing or other protocols and therefore cannot be used to directly initiate standards-track protocol changes.

Existing intra-domain SAV mechanisms have problems in terms of validation accuracy and operational overhead. Current uRPF-based mechanisms derive SAV decisions solely from routing or forwarding

state, which is intended to express reachability rather than authorization of source address usage. More generally, current mechanisms lack authoritative information specifically intended for source address validation that can be consistently and automatically consumed by SAV mechanisms. As a result, uRPF-based mechanisms may not provide accurate validation in scenarios such as asymmetric routing or hidden prefixes (Section 4). Existing ACL-based SAV deployments may have limited applicability in dynamic environments when they rely on operator-driven ACL maintenance. These problems motivate the first two requirements below (in Section 5.1 and Section 5.2). The remaining three requirements (in Section 5.3, Section 5.4, and Section 5.5) are motivated by deployment and operational considerations.

5.1. Accurate Validation

Any new intra-domain SAV mechanism must improve SAV accuracy over existing intra-domain SAV mechanisms. This improvement can be reflected in reduced improper blocks (false positives), reduced improper permits (false negatives), or both. Furthermore, it must seek to mitigate improper blocks and improve the ability to reject spoofed traffic in the gap scenarios described in Section 4. To support this, additional information beyond the local FIB, such as SAV-specific information, may be needed to make validation decisions and generate accurate SAV rules.

5.2. Automatic Updates

Any new intra-domain SAV mechanism must be capable of automatically collecting and processing relevant information, and updating the corresponding SAV rules in response to relevant information changes. 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.

5.3. Incremental Deployment Support

Any new intra-domain SAV mechanism must support incremental deployment and provide measurable benefits even when deployed on only a subset of external interfaces facing hosts or customer networks with no AS.

5.4. No Adverse Impact on Routing Convergence and Fast Reroute

If any new intra-domain SAV mechanism requires disseminating SAV-specific information among intra-domain routers via a protocol, it must not adversely affect the convergence of existing routing protocols or the operation of fast-reroute mechanisms.

5.5. 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.

6. Security Considerations

This document discusses the problems with existing intra-domain SAV practices and identifies informational requirements for new intra-domain SAV mechanisms. As it does not specify any new protocol/mechanism or protocol extension, it does not introduce new security considerations.

7. IANA Considerations

This document does not request any IANA allocations.

8. Acknowledgements

The authors thank Jared Mauch, Joel Halpern, Aijun Wang, Michael Richardson, Gert Doering, Tony Przygienda, Yingzhen Qu, James Guichard, Ron Bonica, Xueyan Song, and others for their valuable comments. The authors also thank Kotikalapudi Sriram for his suggestions on the definition of intra-domain SAV. The authors thank the IETF Directorates and the IESG for their reviews and comments, which helped improve the clarity of this document.

9. References

9.1. 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>>.
- [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>>.

9.2. 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?>>>.

Authors' Addresses

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

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

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

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

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