IDR
Internet-Draft
Intended status: Standards Track
Expires: 22 October 2026

D. Li
Tsinghua University
S. Yue
China Mobile
Z. Li
Huawei Technologies
L. Qin
Zhongguancun Laboratory

20 April 2026

# Source Address Validation at Intra-domain Network Boundary Using BGP
draft-li-idr-savnet-intra-domain-bgp-01

## Abstract

This document proposes a solution for Source Address Validation (SAV) at the intra-domain network boundary using BGP. Routers at the boundary automatically generate accurate SAV rules by using routing information and SAV-specific information. These rules construct a validation boundary which checks the validity of the source address of any data packets flowing into the intra-domain network. This document also introduces BGP extensions for communicating SAV-specific information among routers at the boundary.

## 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 22 October 2026.

## Copyright Notice

Table of Contents



## 1. Introduction

Source address validation (SAV) is essential for mitigating source address spoofing attacks ([RFC6959]) on the Internet. The current operational intra-domain SAV mechanisms include Access List Control (ACL) [RFC2827] and unicast Reverse Path Forwarding [RFC3704]. Their technical problems are detailed in [I-D.ietf-savnet-intra-domain-problem-statement] and a new intra-domain SAV solution that can generate accurate SAV rules in an automatic way is needed.

This document proposes a solution for SAV at the intra-domain network boundary using BGP. We refer to both edge routers and border routers as routers at the boundary of an intra-domain network. Routers at the boundary automatically generate accurate SAV rules by using routing information and SAV-specific information. These rules construct a validation boundary which checks the validity of the source address of any data packets flowing into the intra-domain network. This document also introduces BGP extensions for communicating SAV-specific information among routers at the boundary.

### 1.1. Terminology

SAV Rule: The rule that describes the mapping relationship between a source address (prefix) and its valid incoming router interface(s).

SAV-specific Information: The information specialized for SAV rule generation, which is exchanged among routers.

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.

Edge Router: A router connected to hosts or non-BGP customer networks of the local AS. It forwards traffic from hosts or non-BGP customer networks into the intra-domain network.

Border Router: A router positioned at the boundary between the local AS and one or more external Autonomous Systems (ASes). It runs EBGP and exchanges routing information with external peers.

### 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. SAV Procedure at Edge Routers

Edge routers will generate a prefix allowlist on interfaces facing a non-BGP customer network or a set of host, including only prefixes that can be used as the source address by the non-BGP customer network or the set of hosts. For example, in Figure 1, the prefix allowlist on Intf. 1 should only include all the source prefixes of non-BGP customer network1, the prefix allowlist on Intf. 2 and Intf. 3 should only include all the source prefixes of non-BGP customer

network2, and the prefix allowlist on Intf.4 should only include the network segment of the hosts.

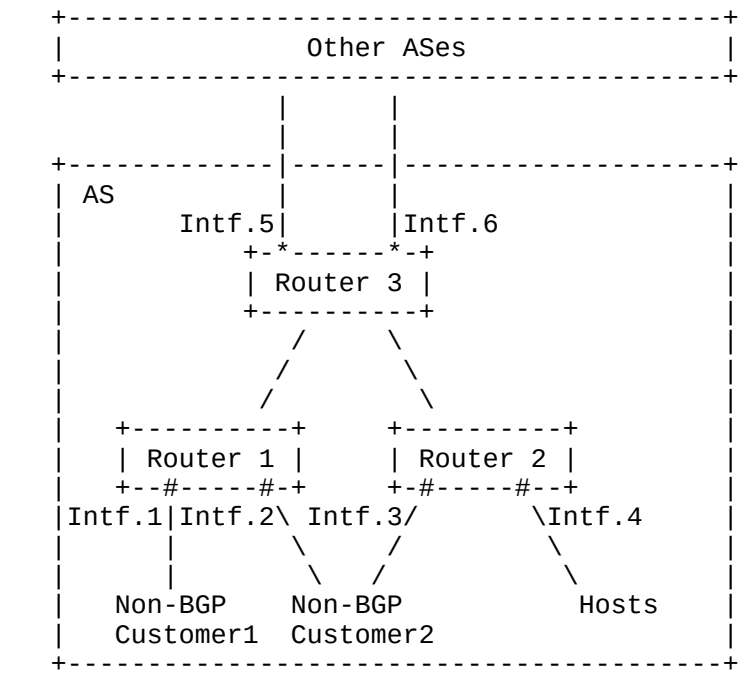


Figure 1: An example of SAV at the boundary of an intra-domain network

To construct a prefix allowlist on the specific interface, the edge router inspects its local Routing Information Base (RIB) and identifies destination prefixes that are reachable via that interface. These prefixes should cover the complete source address space of the connected hosts, single-homed non-BGP customer networks, or multi-homed non-BGP customer networks with symmetric routing.

In asymmetric routing scenarios, these prefixes derived from the local RIB may fail to cover the complete source address space of multi-homed non-BGP customer networks. For example, in Figure 1, the prefix that is reachable via Intf.2 may be part of the complete source address space of non-BGP customer network2 because the other part is only reachable via Intf.3. One approach is to configure static routes with lower priority on the Intf.2 and Intf.3 of Router1 to achieve symmetric routing, without influencing the Forwarding Information Base (FIB) . However, static configuration requires additional manual overhead.

To address this limitation, the Source Prefix Advertisement (SPA) process [I-D.li-savnet-source-prefix-advertisement] is required to enable the construction of a more accurate and complete allowlist. During the SPA process, edge routers will announce these prefixes together with other SAV-specific information. The SAV-specific information includes:

* Multi-homing Interface Group Type (MIIG-Type): It indicates the type of the interface that learns the prefix. In general, there are two types of interfaces: single-homing interface (e.g., Intf.1 and Intf.4 in Figure 1) and multi-homing interface (e.g., Intf.2 and Intf.3 in Figure 1).

* Multi-homing Interface Group Tag (MIIG-Tag): It is to identify the non-BGP customer network of the prefix. Prefixes belonging to the same non-BGP customer network MUST have the identical MIIG-Tag value. Different non-BGP customer networks MUST have different MIIG-tag values.

* (Only) Source Flag: It indicates whether the prefix is an internal-use-only prefix. By default, the flag is set because most prefixes are internal-use-only. For prefixes which are also used to source traffic by other ASes, the flag should be unset.

The detailed procedure for the prefix allowlist generation is as follows:

1. For each interface connecting a non-BGP customer network or a set of hosts, create a set of unique prefixes that are reachable via that interface by inspecting the local RIB. Call it Set A.

2. Considering prefixes and SAV-specific information provided by other edge routers, create a set of unique prefixes that have the same MIIG-Type and MIIG-Tag as prefixes in Set A. Call it Set B.

3. Form the union of Set A and Set B. Apply the union set at the interface as a prefix allowlist.

## 3. SAV Procedure at Border routers

Border routers aim to generate a prefix blocklist on interfaces facing an external AS, including prefixes that can be only used as the source address by the local AS (i.e., internal source addresses). For example, in Figure 1, the prefix blocklist on Intf. 5 and Intf. 6 should include the prefixes of non-BGP customer network1, non-BGP customer network2, and hosts.

The detailed procedure for the prefix blocklist generation is as follows:

1. Considering prefixes and SAV-specific information provided by edge routers, creat a set of unique prefixes with Source Flag being set. Call it Set C.

2. Apply Set C at the interface facing an external AS as a prefix blocklist.

## 4. BGP Extensions for Intra-domain SAVNET

### 4.1. BGP Protocol Relationship

The BGP extensions for SPA communication follow a backward compatible manner without impacting existing BGP functions. New BGP SAVNET subsequent address families will be introduced under the IPv4 address family and the IPv6 address family, respectively. The BGP UPDATE message (specifically the MP_REACH_NLRI and the MP_UNREACH_NLRI attributes) and the BGP Refresh message will be extended. AFI and SAFI will be used for distinguishing the BGP SAVNET messages from other messages.

A few existing path attributes such as Originator_ID and Clister_list or newly defined path attributes MAY be used for BGP SAVNET. Actually, most existing path attributes are not necessarily required for BGP SAVNET. However, if the unnecessary path attributes are carried in BGP updates, they will be accepted, validated, and propagated consistent with the BGP protocol.

### 4.2. Full-mesh IBGP Peering

Edge or border routers enabling BGP SAVNET MUST establish full-mesh iBGP sessions either through direct iBGP sessions or route-reflectors. SAVNET messages within an AS can be advertised through the full-mesh BGP SAVNET sessions. The extensions of BGP messages for carrying SAVNET messages will be introduced in Section 4.3.

### 4.3. BGP SAVNET Protocol Extension

#### 4.3.1. BGP SAVNET SAFI

To make good isolation with existing BGP services, this document defines BGP SAVNET SAFIs under the IPv4 address family and the IPv6 address family, respectively. The values require IANA registration as specified in Section 8. Two BGP SAVNET speakers MUST establish a BGP SAVNET peer and MUST exchange the Multiprotocol Extensions Capability [RFC5492] to ensure that they are both capable of processing BGP SAVNET messages properly.

#### 4.3.2. BGP SAVNET NLRI

The BGP SAVNET NLRI is used to transmit SPA messages (either IPv4 or IPv6). The BGP SAVNET NLRI TLVs are carried in BGP UPDATE messages as (1) route advertisement carried within Multiprotocol Reachable NLRI (MP_REACH_NLRI) [RFC4760], and (2) route withdraw carried within Multiprotocol Unreachable NLRI (MP_UNREACH_NLRI).

While encoding an MP_REACH_NLRI attribute containing BGP SAVNET NLRI TLVs, the "Length of Next Hop Network Address" field SHOULD be set to 0 upon the sender. The "Network Address of Next Hop" field SHOULD not be encoded upon the sender, because it has a 0 length and MUST be ignored upon the receiver.

#### 4.3.3. SPA TLVs

The BGP SAVNET NLRI TLV each carries an SPA message including a source prefix and related information. Therefore, the NLRI TLV is called SPA TLV. This type of TLVs are used in SPA process within an AS. The format is shown below:

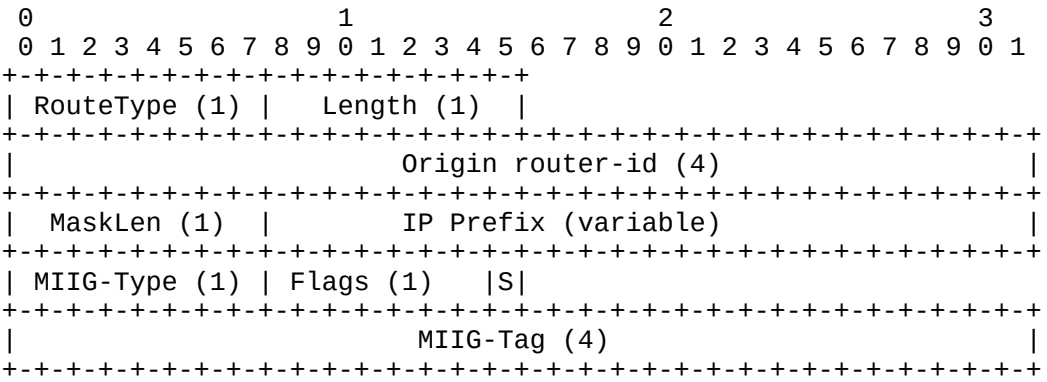


Figure 2: SPA TLV format

The meaning of these fields are as follows:

* RouteType (key): Type of the BGP SAVNET NLRI TLV, the value is 1 for SPA TLV within an AS.

* Length: The length of the BGP SAVNET NLRI value, the RouteType and Length fields are excluded.

* Origin router-id (key): The router ID of the originating node of the source prefix in the deployment domain.

* MaskLen (key): The mask length in bits, which also indicates the valid bits of the IP Prefix field.

* IP Prefix (key): IP address. The length ranges from 1 to 4 bytes for IPv4 and ranges from 1 to 16 bytes for IPv6. Format is consistent with BGP IPv4/IPv6 unicast address.

* MIIG-Type (non-key): Multi-homing Ingress Interface Group Type.

  - Type value 0: Unknown. Indicates that this prefix does not come from any non-BGP customer networks or hosts. It can be a local prefix or a local domain prefix.

  - Type value 1: Single-homing interface. Indicates that this prefix comes from a non-BGP customer network or a set of hosts that is single-homed to the local domain.

  - Type value 2: Multi-homing interface. Indicates that this prefix comes from a non-BGP customer network or a set of hosts that is multi-homed to the local domain, and is connected only to the local domain.

  - Type value 3~255: Reserved for future use.

* Flags (non-key): Bitmap, indicating the attribute flag of the SPA prefix, currently taken:

  - bit 0 (S bit) : Source Flag. The value of 1 indicates that the SPA prefix is single-homing. The value of 0 indicates that the SPA prefix is multi-homing.

* MIIG-Tag (non-key): Multi-homing Ingress Interface Group Tag. The value ranges from 1 to 0xFFFFFFFE. The value 0 is invalid and the value 0xFFFFFFFF is reserved.

## 5. Decision Process with BGP SAVNET

The Decision Process described in [RFC4271] works to determines a degree of preference among routes with the same prefix. The Decision Process involves many BGP Path attributes, which are not necessary for BGP SAVNET SPA process, such as next-hop attributes and IGP-metric attributes. Therefore, this document introduces a simplified Decision Process for SAVNET SAFI.

The purpose of SPA is to maintain a uniform Source Prefix list, which is the mapping from original router-id to IP addresses, across all routers in the deploy domain. To ensure this, it is RECOMMENDED that all routers deploy no ingress or egress route-policies for BGP SAVNET.

### 5.1. BGP SAVNET NLRI Selection

The Decision Process described in [RFC4271] no longer apply, and the Decision Process for BGP SAVNET NLRI are as follows:

1. The locally imported route is preferred over the route received from a peer.

2. The route received from a peer with the numerically larger originator is preferred.

3. The route received from a peer with the numerically larger Peer IP Address is preferred.

#### 5.1.1. Self-Originated NLRI

BGP SAVNET NLRI with origin router-id matching the local router-id is considered self-originated. All locally imported routes should be considered self-originated by default.

Since the origin router-id is part of the NLRI key, it is very unlikely that a self-originated NLRI would be received from a peer. Unless a router-id conflict occurs due to incorrect configuration. In this case, the self-originated NLRI MUST be discarded upon the receiver, and appropriate error logging is RECOMMENDED.

On the other hand, besides the route learn from peers, a BGP SAVNET speaker MUST NOT advertise NLRI which is not self-originated.

### 5.2. BGP Source Prefix Filtering

In actual network deployment, operaters may need to specify whether some source prefixes requires filtering, and BGP community attributes can be used to identify the source prefixes that require filtration.

## 6. Error Handling

### 6.1. Process of BGP SAVNET NLRIs

When a BGP SAVNET speaker receives a BGP Update containing a malformed MP_REACH_NLRI or MP_UNREACH_NLRI, it MUST ignore the received TLV and MUST NOT pass it to other BGP peers. When discarding a malformed TLV, a BGP SAVNET speaker MAY log a specific error.

If duplicate NLRIs exist in a MP_REACH_NLRI or MP_UNREACH_NLRI attribute, only the last one SHOULD be used.

### 6.2. Process of BGP SAVNET SPA TLVs

When a BGP SAVNET speaker receives an SPA TLV with an undefined type, it SHOULD be ignored or stored without parsing.

When a BGP SAVNET speaker receives an SPA TLV with a 0 origin router-id, or the origin router-id is the same as the local router-id, it MUST be considered malformed.

When a BGP SAVNET speaker receives an SPA TLV with an invalid MaskLen field, which is out of the range 1~32 for IPv4 and 1~128 for IPv6, it MUST be considered malformed.

When a BGP SAVNET speaker receives an SPA TLV with an address field, whose length in bytes do not match with the remaining data, it MUST be considered malformed.

When a BGP SAVNET speaker receives an SPA TLV with an unsupported MIIG-Type, it SHOULD be ignored or stored without parsing.

When a BGP SAVNET speaker receives an SPA TLV with a MIIG-Type 0 (Unknown), its MIIG-Tag MUST also be 0, vice versa. Otherwise this SPA TLV MUST be considered malformed.

When a BGP SAVNET speaker receives a malformed SPA TLV, it MUST ignore the received TLV and MUST NOT pass it to other BGP peers. When discarding a malformed TLV, a BGP SAVNET speaker MAY log a specific error.

When a BGP SAVNET speaker processes Flags in an SPA TLV, the defined bits MUST be processed and the undefined bits MUST be ignored.

## 7. Security Considerations

TBD.

## 8. IANA Considerations

The BGP SAVNET SAFIs under the IPv4 address family and the IPv6 address family need to be allocated by IANA.

## Acknowledgements

TBD.

## References

### Normative References

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <https://www.rfc-editor.org/info/rfc4760>.

[RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <https://www.rfc-editor.org/info/rfc5492>.

[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>.

### 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/info/rfc4271>.

[RFC6959] McPherson, D., Baker, F., and J. Halpern, "Source Address Validation Improvement (SAVI) Threat Scope", RFC 6959, DOI 10.17487/RFC6959, May 2013, <https://www.rfc-editor.org/info/rfc6959>.

[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>.

[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>.

[I-D.ietf-savnet-intra-domain-problem-statement] Li, D., Wu, J., Qin, L., Huang, M., and N. Geng, "Source Address Validation in Intra-domain Networks Gap Analysis, Problem Statement, and Requirements", Work in Progress, Internet-Draft, draft-ietf-savnet-intra-domain-problem-statement-23, 14 April 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-23>.

[I-D.li-savnet-source-prefix-advertisement] Qin, L., Geng, N., and D. Li, "Source Prefix Advertisement for Intra-domain SAVNET", Work in Progress, Internet-Draft, draft-li-savnet-source-prefix-advertisement-06, 29 January 2026, <https://datatracker.ietf.org/doc/html/draft-li-savnet-source-prefix-advertisement-06>.

## Authors' Addresses

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

Shengnan Yue
China Mobile
Beijing
China
Email: yueshengnan@chinamobile.com

Zhenbin Li
Huawei Technologies
Beijing
China
Email: lizhenbin@huawei.com

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