

SAVNET  
Internet-Draft  
Intended status: Standards Track  
Expires: December 30, 2025

W. Cheng  
China Mobile  
C. Lin  
New H3C Technologies  
S. Yue  
China Mobile  
June 30, 2025

Intra-domain SAV Support via BGP  
draft-cheng-savnet-intra-domain-sav-bgp-03

## Abstract

This document describes a method for publishing source prefixes via the BGP protocol, iterating through the SAV table entries based on intra-domain next hop SAV rules. The generation of intra-domain next hop SAV rules is implemented by the intra-domain IGP protocol, and the BGP protocol inherits the source interface list from its next hop SAV rules to generate the SAV rule table for source prefixes.

## 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 December 30, 2025.

## 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
2. Terminology.....	3
3. Solution.....	3
3.1. Overview.....	3
3.2. Intra-domain SAVNET.....	6
3.3. Procedure.....	8
3.4. Multi-Homed Scenarios.....	8
3.5. BGP Source Prefix Filtering.....	10
4. Example.....	10
5. Deployment Considerations.....	11
5.1. Intra-domain SAV Support via IGP and BGP.....	11
6. IANA Considerations.....	11
7. Security Considerations.....	12
8. References.....	12
8.1. Normative References.....	12
8.2. Informative References.....	12
Acknowledgments.....	12
Authors' Addresses.....	12

## 1. Introduction

As shown in Figure 1, the existing network has the following scenario: within the intra-domain network, topology information is disseminated via the IGP protocol, while prefix information is distributed via iBGP neighbors. All iBGP nodes establish iBGP neighbor relationships with Route Reflectors (RRs) and exchange source prefix information. The IGP protocol contains network topology information but lacks source prefix information, while the BGP protocol holds source prefix information but does not include network topology information. In this scenario, it is necessary to combine the network topology information from IGP with the source prefix information from the BGP protocol in order to compute the source prefix's associated source port information and generate SAV rules.






Figure 1: The case of the SAVNET Procedure

The scenario described in [draft-cheng-savnet-intra-domain-sav-IGP-00] and [I-D.lin-Intra-domain-savnet-method] involves the publication of SAV source prefix by the IGP protocol, and the generation of SAV rules based on the connectivity calculation using the IGP's topology information. However, for the scenario described in this document, where the source prefix information is published by the BGP protocol, it is unable to generate the required SAV rules.

This document describes how to generate SAV rules using the topology information from the IGP protocol and the source prefix information from the BGP protocol in this network scenario.

## 2. Terminology

The following terminologies are used in this document.

**SAV Rule:** The rule that indicates the source validity of a specific IP address or an IP prefix.

**SAV Table:** The table or data structure that implements the SAV rules and is used for source address validation in the data plane.

**IGP:** Interior Gateway Protocol.

**BGP:** Border Gateway Protocol.

**Source prefix:** The source prefixes are used to validate source addresses in the data plane.

## 3. Solution

### 3.1. Overview

As shown in Figure , [draft-ietf-savnet-intra-domain-architecture-00] describes SAV filtering on three types of interfaces: interfaces

on host-facing routers, such as interface 1 on Router C as shown in the figure; interfaces on customer-facing routers, such as interface 1 on Router A and interface 1 on Router B as shown in the figure; and interfaces on AS border routers, such as interface 2 on Router D and interface 2 on Router E when facing another AS.

For the first two types of interfaces, SAV filtering is primarily carried out using an allow-list. For the third type of interface, filtering is primarily implemented using a block-list.

This document also describes SAV filtering on the interfaces of routers within the domain, such as Interface 1 of Router D and Interface 1 of Router E, as depicted in the diagram. Such interfaces employ allow-list based SAV filtering, with the computation process outlined in Section 3.3.

The specific details about allow-list and block-list are detailed below.

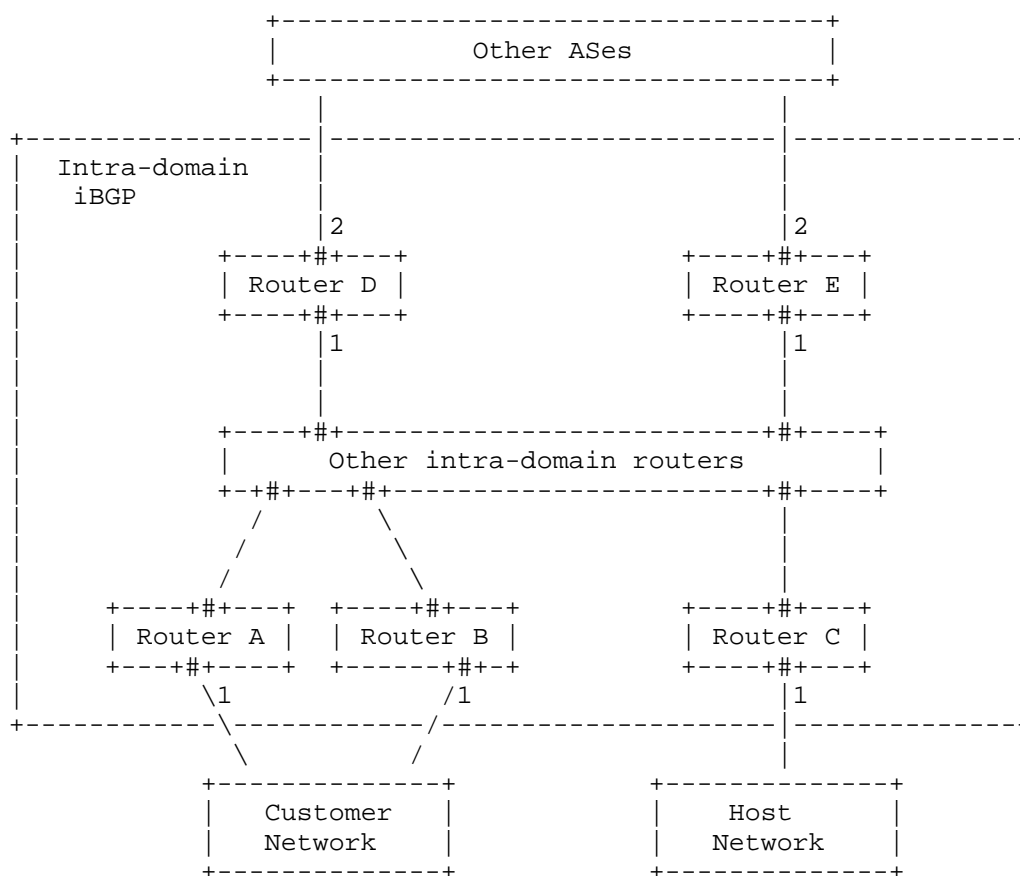


Figure 2: Overview of intra-domain SAVNET architecture

The term allow-list means that if a packet received from a specified interface has a source address found in the interface's allow-list, then the packet should be permitted to pass through. For example, for an allow-list entry (P1, A-1), if a packet received from interface A-1 has a source address within the range of P1, then the packet will be allowed to pass through. Otherwise, the packet will be discarded.

The term block-list means that if a packet received from a specified interface has a source address found in the interface's block-list, then the packet must be discarded. Otherwise, the packet will be allowed to pass through.

For the host-facing router interfaces, the allow-list can be generated for directly connected host segments, as shown in Figure 2. If the interface 1 network segment address of Router C is P, the allow-list (P, C-1) can be generated for Router C's interface 1.

For the ports of routers facing the customers, as shown in Figure 2, if the interface 1 network segment address of Router A is P1, and the interface 1 network segment address of Router B is P2, the allow-lists (P1, A-1) and (P2, A-1) can be generated for Router A's interface 1, and the allow-lists (P1, B-1) and (P2, B-1) can be generated for Router B's interface 1.

For the router interfaces facing the customers, as shown in Figure 2, where Router A's interface 1 belongs to network segment address P1 and Router B's interface 1 belongs to network segment address P2, the allow-lists (P1, A-1) and (P2, A-1) can be generated for Router A's interface 1, and the allow-lists (P1, B-1) and (P2, B-1) can be generated for Router B's interface 1.

For the inter-domain router interfaces, as depicted in Figure 4, where Router D's interface 1 and Router E's interface 1 are involved, the calculation process described in section 4.3 can be used to generate the allow-lists for Router D's interface 1 and Router E's interface 1.

For the interface facing another AS, as shown in Figure 2, such as Router D's interface 2 and Router E's interface 2, this scenario typically involves using block-lists for SAV filtering. The generation of block-lists on the edge interfaces connecting to external domains can be achieved by generating block-list entries on the upper interface when learning routes from the lower interface.

For example, if Router D learns the prefix P on interface 1, then a block-list entry (P, D-2) can be generated on interface 2.

### 3.2. Intra-domain SAVNET

This section introduces a new method for computing and generating SAV rules based on BGP source prefix and IGP topology information in an intra-domain scenario. This method relies on two fundamental pieces of information: the source prefix information and reachability information. The source prefix information can be transmitted through static configuration or the BGP protocol. This document addresses the scenario where source prefix information is transmitted via the BGP protocol.

The source prefix information consists of the source prefix and the next-hop information for the prefix publication.

The IGP's topology information includes the connectivity details between nodes and the IGP prefix information published by each node.

As depicted in Figure 3, source prefix information is disseminated via the BGP protocol, where Router C advertises the source prefix Prefix1 with a next hop of Router 2, Router D advertises the source prefix Prefix2 with a next hop of Router 3, and Router E advertises the source prefix Prefix3 with a next hop of Router 4. Router B, serving as the BGP Route Reflector, is responsible for collecting and reflecting all BGP source prefix information.

Based on the IGP's topology information, the interface list corresponding to the IGP prefix can be calculated. The specific calculation process can be found in [draft-cheng-savnet-intra-domain-sav-BGP-00]. The first-level NextHop SAV rule table is generated based on this information in the form of (IGP-Prefix, if).

The calculation of the next-hop SAV rule is not limited to IGP and can involve other new extended protocols not described in this document.

Subsequently, using the source prefix information distributed via the iBGP protocol, a match is made against the first-level SAV rule table based on the source prefix information. Once a match is found, the interface list "if" is inherited to produce the second-level SAV rule table (BGP-Prefix, if). This document relies on the generation of SAV rules based on the next hop derived from the IGP protocol. The relationships of the generated SAV rule table are illustrated in Figure 4.

+-----+

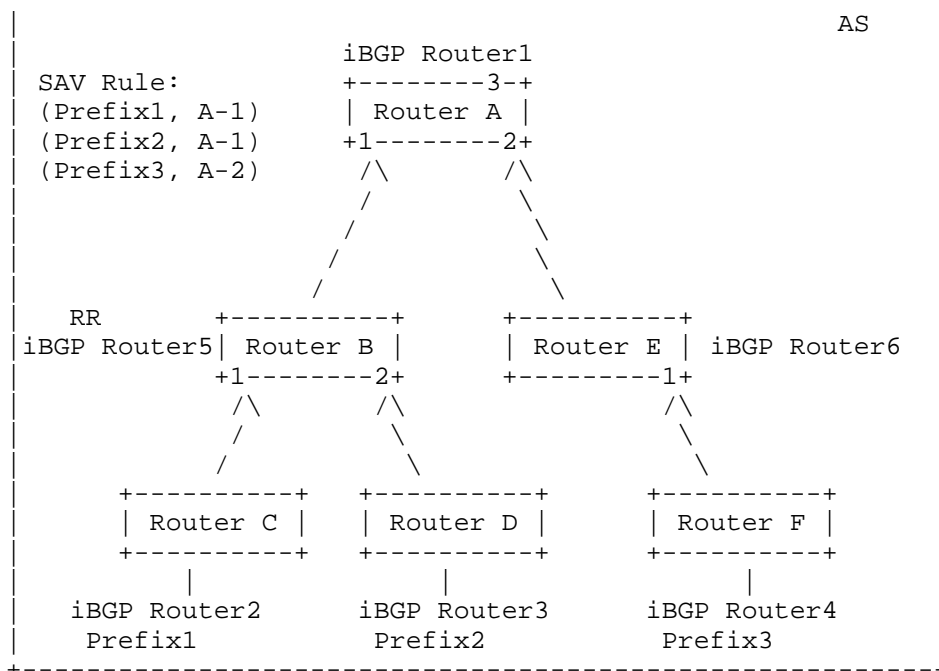
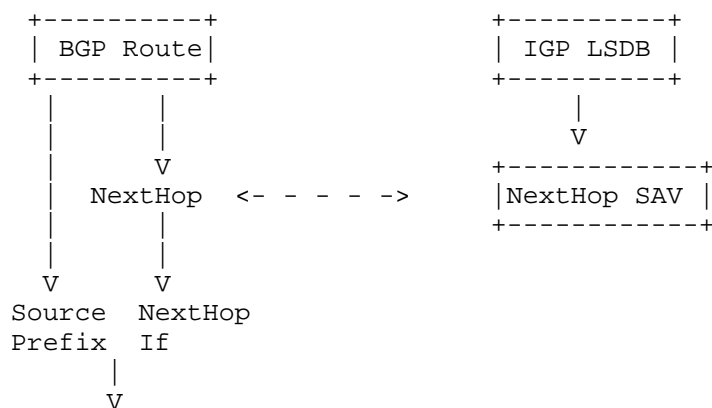


Figure 3: Example 1 of Topology Calculation

This approach enables automatic adjustment of SAV table entries based on topological changes, thereby achieving secure protection for source addresses within the domain.



```

+-----+
| BGP SAV |
+-----+

```

Figure 4: The calculation process of BGP SAV table

### 3.3. Procedure

The calculation process for intra-domain SAV rules based on BGP is as follows:

Step 1: Perform calculation based on the LSDB of the IGP protocol and generate first-level SAV rules using the prefix information published by the IGP nodes. The generated SAV rule takes the form: (IGP-Prefix, if). This also forms the next-hop SAV table required for BGP.

Step 2: Iterate through all source prefix information distributed by the BGP protocol. For each source prefix, match it with the corresponding next-hop information of the publisher. Then, search and match this next-hop address in the SAV rules generated in Step 1. Obtain and utilize the inherited interface list from the first-level SAV rules to generate second-level SAV rules. The generated rules take the form: (BGP-Prefix, if).

Step 3: If there are changes in the topological information of the IGP protocol, repeat the calculation in Step 1. If there are changes in the SAV rules generated in Step 1, the BGP protocol refreshes the (BGP-Prefix, if) rule table based on the next-hop associated SAV table, thus skipping Step 2.

Step 4: If there are changes in the source prefix information distributed by the BGP protocol, skip Step 1 and proceed with the calculation according to Step 2, refreshing the rule list generated by BGP.

### 3.4. Multi-Homed Scenarios

In the multi-homed scenario depicted in Figure 5, the prefix 10.0.0.0/16 is accessed via nodes D and E, with the subnet 10.0.1.0/24 being accessed via node D, and 10.0.0.0/24 being accessed via node E. Through IGP connectivity calculation, the SAV table entries derived on node A are: (10.0.1.0/24, A-1), (10.0.0.0/24, A-2).

```

+-----+
| iBGP 11.11.11.11      AS |
| Node A SAV Rule:  +-----+ |

```






Figure 5: The case of Multi-homing Access

In the scenario of single-homed user network access to the local network, ACL or strict uRPF can generally provide effective protection. However, in the case of multi-homed user network access, there may be asymmetric routing and the prefixes of asymmetric routes on the router may dynamically change to achieve dynamic load balancing. In this case, ACL or strict uRPF cannot be directly used to protect the user access interface as a allow-list, as this may lead to issues of misfiltering.

In the mentioned network, Node C could receive packets with a source address of 10.0.0.0/24 via interface C-1, and Node A could receive similar packets via interface A-1, potentially causing misfiltering. Similarly, Node E could receive packets with a source address of 10.0.1.0/24 via interface E-1, and Node A could receive similar packets via interface A-2, leading to potential misfiltering of packets.

To address this issue , one approach is to configure a static route (10.0.0.0/24, C-1) on the interface C-1 of Node C, and when advertising this route through BGP, specify a higher Prefer value than the Prefer value of this route on Node E. Similarly, on Node E, configure a static route (10.0.1.0/24, E-1) and when advertising it through BGP, specify a higher Prefer value than the Prefer value of this route on Node C.

The resulting SAV entries are as follows:

```
Entries on Node C: (10.0.1.0/24, C-1), (10.0.0.0/24, C-1)
Entries on Node E: (10.0.0.0/24, E-1), (10.1.0.0/24, E-1)
Entries on Node A: (10.0.1.0/24, A-1), (10.0.0.0/24, A-1)
                  (10.0.1.0/24, A-2), (10.0.0.0/24, A-2)
```

3.5. BGP Source Prefix Filtering

In actual network deployment, some source prefixes do not need to be filtered. To distinguish which source prefixes require filtering, BGP community attributes can be used to identify the source prefixes that require filtration.

When the BGP protocol advertises a source prefix, it includes the specified community attribute to indicate that this source prefix requires SAV calculation. On the devices performing SAV calculation, SAV rules are generated only for routes with the specified community attribute.

4. Example

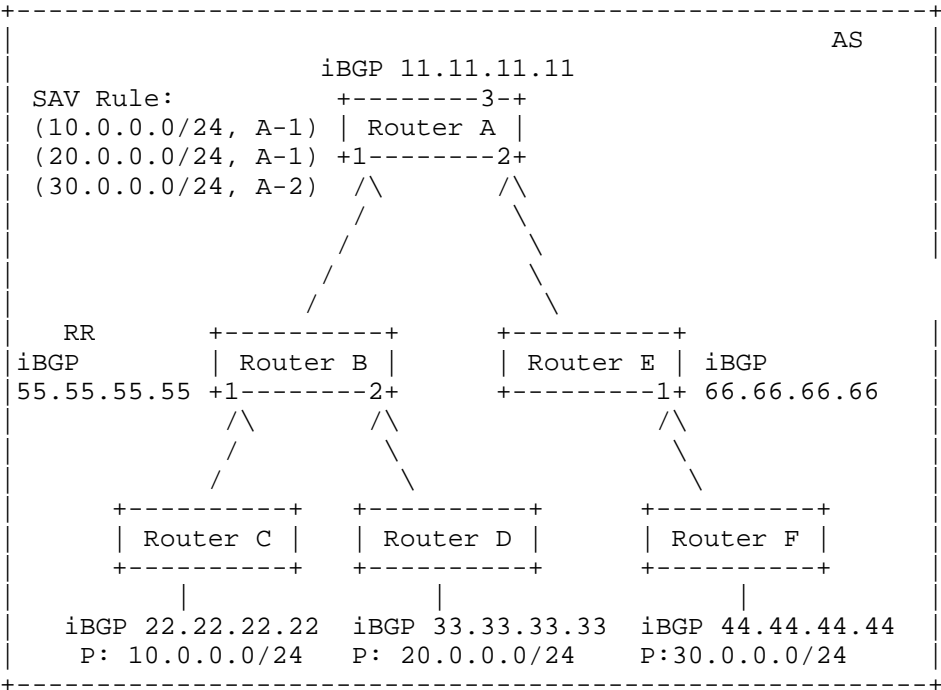


Figure 7: Example 3 of Topology Calculation

After conducting internal IGP calculations, on router A, it is determined that BGP neighbors reachable via A-1 are 22.22.22.22 and 33.33.33.33. BGP neighbor 44.44.44.44 is reachable via A-2.

Following the source prefix calculation in BGP, inheriting the outgoing interface information from the connectivity calculation, router A can compute the following sav table entries: (10.0.0.0/24, A-1) obtained from BGP neighbor 22.22.22.22, (20.0.0.0/24, A-1) obtained from BGP neighbor 33.33.33.33, and (30.0.0.0/24, A-2) obtained from BGP neighbor 44.44.44.44.

## 5. Deployment Considerations

### 5.1. Intra-domain SAV Support via IGP and BGP

If the network topology information and source prefix information within the domain are both conveyed by the IGP protocol, SAV rules can be automatically generated following the calculation method described in [draft-cheng-savnet-intra-domain-sav-IGP-00] or [I-D.lin-Intra-domain-savnet-method].

If in the network, the intra-domain network topology information is conveyed by the IGP protocol, while the intra-domain source prefix information is transmitted via the BGP protocol, this SAV calculation method can be deployed to generate SAV rules for preventing source address attacks in outbound and inbound traffic.

If the intra-domain source prefixes are transmitted via BGP, while network connectivity information is conveyed by protocols other than IGP, this deployment can still be used to calculate SAV rules. The BGP protocol simply inherits the interfaces from the topological calculation into the final generated SAV rules, based on the next-hop information in the source prefixes.

Furthermore, it is also possible to plan a separate BGP domain within the intra-domain, using BGP RR to reflect and propagate all intra-domain source prefixes. First, through IGP or other extended technologies, the SAV table entries corresponding to the next hops of BGP source prefixes are calculated. Finally, through the next hop of BGP, the SAV table entries of the next hops are obtained to generate the BGP-published source prefix SAV table entries, ultimately achieving BGP calculation SAVNET functionality within the intra-domain.

## 6. IANA Considerations

This document does not involve IANA.

## 7. Security Considerations

TBD

## 8. References

### 8.1. Normative References

[I-D.ietf-savnet-intra-domain-architecture]

Li, D., Wu, J., Huang, M., Chen, L., Geng, N., Qin, L.,  
and F. Gao, "Intra-domain Source Address Validation  
(SAVNET) Architecture", Work in Progress, Internet-Draft,  
draft-li-savnet-intra-domain-architecture-03, 25 July  
2023, <<https://datatracker.ietf.org/doc/html/draft-li-savnet-intra-domain-architecture-06>>.

[I-D.lin-Intra-domain-savnet-method] D. Li, "Intra-domain SAVNET  
method", Work in Progress,  
<<https://www.ietf.org/archive/id/draft-lin-savnet-lsr-intra-domain-method-03.txt>>

### 8.2. Informative References

[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-02, 17 August 2023,  
<<https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-02>>.

## Acknowledgments

TBD

## Authors' Addresses

Weiqiang Cheng  
China Mobile  
China

Email: [chengweiqiang@chinamobile.com](mailto:chengweiqiang@chinamobile.com)

Changwang Lin  
New H3C Technologies  
China

Email: [linchangwang.04414@h3c.com](mailto:linchangwang.04414@h3c.com)

Shengnan Yue  
China Mobile  
China  
[yueshengnan@chinamobile.com](mailto:yueshengnan@chinamobile.com)

