

IDR
Internet-Draft
Intended status: Standards Track
Expires: 25 September 2026

N. Geng
Z. Li
Z. Tan
M. Liu
Huawei Technologies
D. Li
Tsinghua University
24 March 2026

BGP Extensions for Source Address Validation Networks (BGP SAVNET)
draft-geng-idr-bgp-savnet-06

Abstract

Many source address validation (SAV) mechanisms have been proposed for preventing source address spoofing. However, existing SAV mechanisms are faced with the problems of inaccurate validation or high operational overhead in some scenarios. This document proposes BGP SAVNET by extending BGP protocol for SAV. This protocol can propagate SAV-related information through BGP messages. The propagated information will help edge/border routers automatically generate accurate SAV rules. These rules construct a validation boundary for the network and help check the validity of source addresses of arrival data packets.

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 25 September 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	3
1.1. Terminology	4
1.2. Requirements Language	5
2. BGP Protocol Relationship	5
3. BGP SAVNET Solution	5
3.1. Goals	6
3.2. Intra-domain BGP SAVNET	7
3.3. Inter-domain BGP SAVNET	10
4. BGP SAVNET Peering Models	12
4.1. Full-mesh IBGP Peering	13
4.2. EBGP Peering between ASes	13
5. BGP SAVNET Protocol Extension	13
5.1. BGP SAVNET SAFI	13
5.2. BGP SAVNET NLRI	13
5.2.1. SPA TLVs for Intra-domain BGP SAVNET	14
5.2.2. SPA TLVs for Inter-domain BGP SAVNET	15
5.3. BGP SAVNET Refresh	16
5.3.1. The SPD TLVs for Inter-domain BGP SAVNET	17
6. Decision Process with BGP SAVNET	18
6.1. BGP SAVNET NLRI Selection	18
6.1.1. Self-Originated NLRI	19
7. Error Handling	19
7.1. Process of BGP SAVNET NLRIs	19
7.2. Process of BGP SAVNET SPA TLVs	19
7.3. Process of BGP SAVNET Refresh	20
7.4. Process of BGP SAVNET SPD TLVs	20
8. Convergence Considerations	21
9. Deployment Considerations	21
10. Security Considerations	21
11. IANA Considerations	21
Acknowledgements	22
References	22
Normative References	22
Informative References	23
Authors' Addresses	24

1. Introduction

Source address validation (SAV) is essential for preventing source address spoofing attacks ([RFC6959]) and helps trace back network attackers. For a network, SAV mechanisms can be deployed on edge routers or border routers for validating the packets from the connected subnets or neighboring ASes [manrs-antispoofing].

ACL-based ingress filtering ([RFC2827], [RFC3704]) and Source-based RTBH ([RFC5635]) can be used for source address filtering. However, the two mechanisms are not specific for SAV. High operational overhead may be induced if they are managed mostly by manual configurations

[I-D.ietf-savnet-intra-domain-problem-statement][I-D.ietf-savnet-inter-domain-problem-statement].

Many SAV mechanisms, such as strict uRPF, loose uRPF, FP-uRPF, VRF-uRPF, and EFP-uRPF ([RFC3704], [RFC8704]), leverage local routing information (FIB/RIB) to automatically generate SAV rules. The rules indicate the wanted incoming interfaces of source addresses and deny source addresses coming from unwanted interfaces

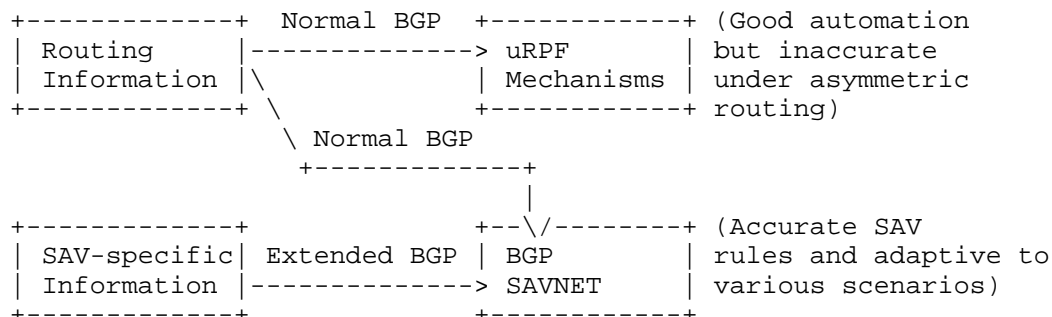
[I-D.huang-savnet-sav-table]. The uRPF mechanisms can achieve good automation but may have inaccurate validation problems under asymmetric routing

[I-D.ietf-savnet-intra-domain-problem-statement][I-D.ietf-savnet-inter-domain-problem-statement].

This is because these uRPF mechanisms are "single-point" designs. They leverage the local FIB or local RIB table to determine the valid incoming interfaces for source addresses, which may not match the real incoming interfaces. That is, purely relying on local routing information for SAV is not enough for achieving both good automation and high accuracy

This document proposes extensions of BGP protocol for SAV networks, named as BGP SAVNET. Unlike existing "single-point" mechanisms, BGP SAVNET allows coordination between the routers within a network or in different ASes by propagating SAV-specific information through extended BGP messages

[I-D.li-savnet-intra-domain-architecture][I-D.wu-savnet-inter-domain-architecture]. SAV-specific information supplements the missing part of the local route information and assists routers to generate accurate SAV rules. The following figure shows a comparison of existing uRPF mechanisms and BGP SAVNET.



The BGP SAVNET protocol is suitable to generating SAV rules for both IPv4 and IPv6 addresses. The SAV rules can be used for validating any native IP packets or IP-encapsulated packets.

1.1. Terminology

SAV: Source address validation, an approach to preventing source address spoofing.

SAV Rule: The rule that indicates the valid incoming interfaces for a specific source prefix.

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

Internal (or Local) Source Address: The source addresses owned by the subnets of local AS. The source addresses of the connection link between subnets and local AS can also be considered as internal source addresses.

External (or Remote) Source Address: The source addresses owned by other ASes. Some source addresses like anycast addresses can be both internal and external source addresses.

Local Routing Information: The information in a router's local RIB or FIB that can be used to infer SAV rules.

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

Edge Router: An intra-domain router for an AS that is directly connected to a subnet of the AS.

Border Router: An intra-domain router for an AS that is connected to other ASes. A router in an AS can be both an edge router and a border router, if it is connected to both the AS's subnets and other ASes.

Source AS: The AS whose source prefixes need to be validated at Validation AS.

Validation AS: The AS that conducts SAV for the source prefixes of Source AS.

SPA: Source prefix advertisement, i.e., the process for advertising the origin source addresses/prefixes of a router or an AS.

SPD: Source path discovery, i.e., the process for discovering the real incoming directions of particular source addresses/prefixes.

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. BGP Protocol Relationship

The BGP extensions for BGP SAVNET 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 Cluster_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.

3. BGP SAVNET Solution

3.1. Goals

For an AS, the goal of BGP SAVNET is to construct a validation boundary for the AS. SAV-specific information propagated by extended BGP messages can assist edge and border routers on the network boundary to generate SAV rules. Edge routers connected to subnets generate rules for validating packets from users, while border routers connected to other ASes generate rules for validating packets from other ASes. Figure 1 shows the example of validation boundary for an AS

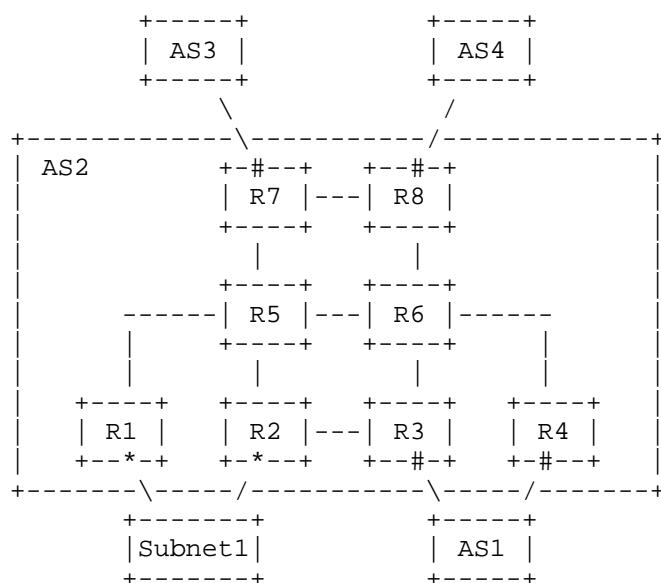


Figure 1: An example of validation boundary for an AS

From a perspective of an AS, source addresses can be largely classified into two categories: internal (or local) source address and external (or remote) source address. The BGP SAVNET solution consists two parts: intra-domain BGP SAVNET and inter-domain BGP SAVNET. The parts of solution focus on the validation of internal and external source address, respectively.

- * Intra-domain BGP SAVNET: SAV for protecting internal source addresses. In Figure 1, it can be deployed at '*' or '#' to restrict a subnet to use only its own internal source addresses and to block external packets from other ASes with any internal source addresses. SAV rules are generated without any cooperation or interactions (such as prefix advertisements) between the local AS and subnets/other ASes.

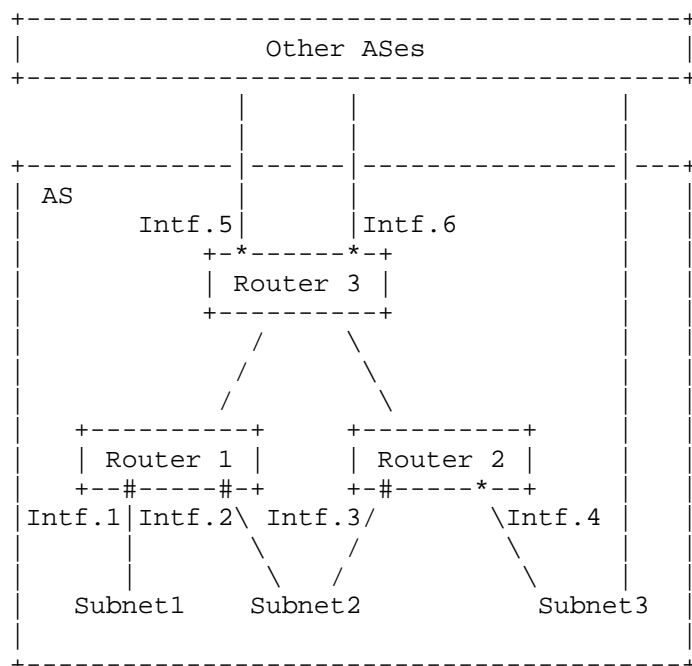
- * Inter-domain BGP SAVNET: SAV for protecting external source addresses. In Figure 1, it can be deployed at '#' for blocking the source addresses of packets coming from unwanted directions (i.e., coming from unwanted neighbor ASes). Cooperation or interactions between the local AS and other ASes are required.

3.2. Intra-domain BGP SAVNET

Figure 2 shows an example of intra-domain BGP SAVNET. Router 1 and Router 2 are edge routers that enable SAV at the interfaces connected to subnets. Router 3 is a border router that conducts SAV at the interfaces connected to other ASes.

In general, there are four types of interfaces:

- * Single-homing interface. When a subnet has only one uplink connected to the upper-layer network, the connected interface at the edge router of upper-layer network can be defined as a "Single-homing interface", e.g., Intf.1 in Figure 2.
- * Complete multi-homing interface. When a subnet has dual or multiple uplinks that connect to a single upper-layer network with BGP SAVNET deployed, the connected interfaces at the edge routers of upper-layer network are called "Complete multi-homing interfaces", which corresponds to Intf.2 and Intf.3 in Figure 2.
- * Incomplete multi-homing interface. When a subnet has dual or multiple uplinks that are connected to multiple upper-layer networks, the interfaces at the edge routers of upper-layer network are called the "Incomplete multi-homing interfaces", which corresponds to Intf.4 in Figure 2.
- * Internet interface. It's the external interfaces that are connected to the Internet on border routers. Typically, a network usually has multiple Internet interfaces for load balancing or backup, which corresponds to Intf.5 and Intf.6 in Figure 2.



Intf '#' enables prefix allowlist
 Intf '*' enables prefix blocklist

Figure 2: An example of intra-domain BGP SAVNET

The goal of intra-domain BGP SAVNET is to generate source prefix allowlist or blocklist for the interfaces on edge or border routers. For the "Single-homing interface" and "Complete multi-homing interface", prefix allowlist is applied (i.e., "Interface-based prefix allowlist" mode in [I-D.huang-savnet-sav-table]). The prefix allowlist of an interface should only include all the source prefixes of the connected subnet and denys any source addresses not covered by the prefixes in the list. In Figure 2, the prefix allowlist of Intf. 1 should only include all the source prefixes of Subnet1, and the prefix allowlists of Intf. 2 and Intf. 3 should only include all the source prefixes of Subnet2.

For "Incomplete multi-homing interface" and "Internet interface", prefix blocklist is enabled (i.e., "Interface-based prefix blocklist" mode in [I-D.huang-savnet-sav-table]). For a specific interface, its prefix blocklist should include the internal prefixes that are owned by the subnets connected to "Single-homing interfaces" and "Complete multi-homing interfaces". In Figure 2, the prefix blocklist of Intf. 4, Intf. 5, and Intf. 6 should include all the source prefixes of

Subnet1 and Subnet2. The reason why "Incomplete multi-homing interface" like Intf. 4 not using prefix allowlist is that the local AS itself can hardly learn the complete set of source prefixes of Subnet3 if the subnet advertises asymmetric prefixes to the multi-homed up-layer networks (i.e., the local AS is one of the up-layer networks).

The above goal should be achieved while meeting two requirements of high accuracy (even under asymmetric routing) and good automation. To this end, Source Prefix Advertisement (SPA) process is designed in intra-domain BGP SAVNET solution. During the SPA process, edge routers will proactively announce all the source prefixes learned by local "Single-homing interfaces" and local "Complete multi-homing interfaces" from the connected subnets via SPA messages. Some related information of each announced source prefix will also be propagated together with the source prefix. The related information of each announced source prefix can be:

- * Multi-homing Interface Group Type (MIIG-Type): It indicates the type of the interface that learns the prefix. MIIG-Type MUST be one of the four types mentioned above.
- * Multi-homing Interface Group Tag (MIIG-Tag): It is to identify the subnet of the prefix. The prefixes belonging to the same subnet MUST have the identical MIIG-Tag value. Different subnets MUST have different MIIG-tag values.
- * (Only) Source Flag: It indicates whether the prefix is owned by one subnet. By default, the flag is set because most of the prefixes are owned by one network. For anycast addresses/prefixes or direct server return (DSR) addresses/prefixes [I-D.ietf-savnet-inter-domain-problem-statement], the flag should be unset (possibly manually).

It can be seen that the SPA message of a source prefix includes four parts: source prefix, MIIG-Type, MIIG-Tag, and Source Flag. Next, how to use the SPA messages to generate SAV rules will be introduced.

- * In the case of "Single-homing interface", the prefix allowlist can be generated only through local routing information (i.e., local RIB), without the engagement of SPA messages. The method building the allowlist is, each Dest Prefix in RIB that records this interface as an outgoing interface will be added to the prefix allowlist.
- * In the case of "Complete multi-homing interface", in addition to collecting prefixes of the target interface itself in local RIB, routers also need merge prefixes from the received SPA messages

and other local interfaces into the allowlist to construct a complete list. First, the prefixes in received SPA, which take the same "MIIG-Type" and "MIIG-Tag" values as the target interface, are added to the allowlist. Second, if there are local interfaces having the same "MIIG-Type" and "MIIG-Tag" values, they will share prefixes collected from local RIB into each other's allowlist.

- * Routers with "Incomplete multi-homing interface" or "Internet interface" will generate prefix blocklist for the target interface. First, the prefixes of local "Single-homing interfaces" or "Complete Multi-homing interfaces" on the local router will be put into the blocklist. Second, the prefixes in the received SPA messages which have the MIIG-Types of either "Single-homing interface" or "Complete Multi-homing interface" but with Source Flag being set, will also be added into the blocklist. The prefix with Source Flag being unset will not be included into the blocklist because the prefix is multi-source and the "Incomplete multi-homing interface" or "Internet interface" may be the legitimate incoming interface of the multi-source prefix.

Note that, intra-domain BGP SAVNET solution can also work if the subnet is a stub AS (e.g., the subnets are replaced with stub ASes in Figure 2). The source prefixes of the stub AS can be considered as the internal prefixes of the local AS when conducting the solution.

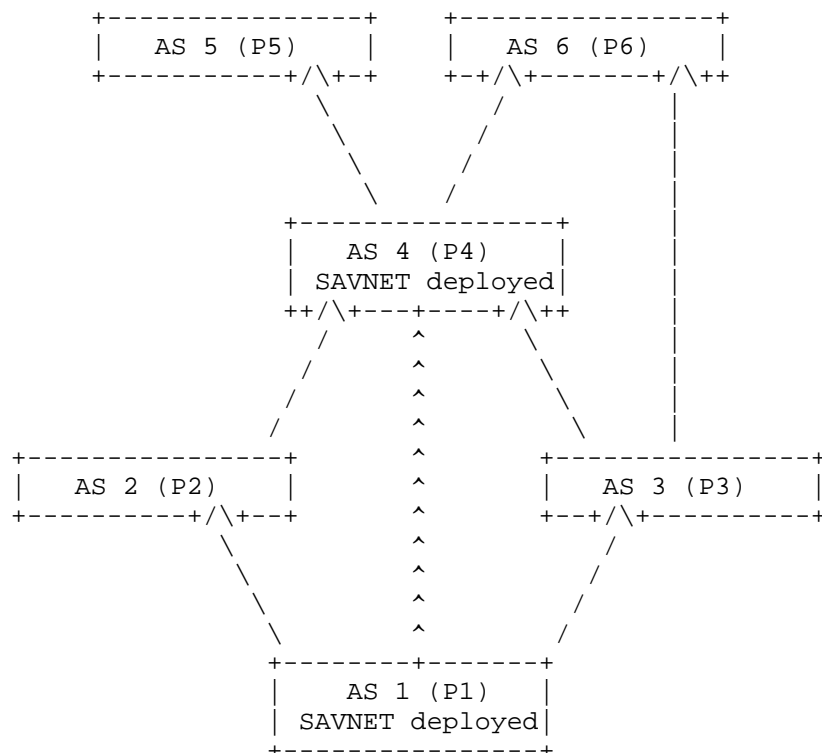
3.3. Inter-domain BGP SAVNET

As described previously, inter-domain BGP SAVNET is for protecting external source addresses that are owned by other ASes (usually remote ASes). Cooperation or interactions between the local AS and other ASes are required.

The local AS that deploys inter-domain BGP SAVNET and conducts validation on the external source addresses is defined as Validation AS. Source AS is defined as the AS whose source prefixes need to be validated at Validation AS. For any AS, it can be configured as Source AS or Validation AS, or it can also act as both Source AS and Validation AS.

The goal of inter-domain BGP SAVNET is to help Validation AS generate prefix blocklist for the source prefixes of Source AS at the proper external interfaces of Validation AS. Which source prefixes that need to be validated and which external interfaces should block these prefixes depend on the indication of Source AS. Inter-domain BGP SAVNET provides the communication channel for Source AS and Validation AS.

Figure 3 shows an example of inter-domain BGP SAVNET. AS 1 and AS 4 have deployed SAVNET on the border routers (i.e., ASBRs) connected to other ASes. Suppose AS 1 is configured as Source AS and AS 4 acts as Validation AS. In the example, AS 4 can help block P1 of AS 1 at the interfaces connected to specific neighbor ASes.



RIB in AS 1:

To P2, preferred AS_PATH = [AS 2]
 To P3, preferred AS_PATH = [AS 3]
 To P4, preferred AS_PATH = [AS 2, AS 4]
 To P5, preferred AS_PATH = [AS 3, AS 4, AS 5]
 To P6, preferred AS_PATH = [AS 3, AS 4, AS 6]

Figure 3: An example of inter-domain BGP SAVNET

When Source AS and Validation AS enable BGP SAVNET, a BGP SAVNET session between the two ASes will be established. Figure 3 shows the session between AS 1 and AS 4 by ">>>>". The solution of inter-domain BGP SAVNET consists of two processes: SPA and Source Path Discovery (SPD).

- * SPA process: Source AS advertises its own AS number and its own source prefixes to Validation AS through SPA messages. SPA messages contain the complete set of source prefixes of Source AS or only the source prefixes that want to be protected. Some hidden source prefixes that do not appear can also be advertised to Validation AS through SPA messages. The advertised source prefixes MUST be authorized to Source AS by RPKI ROAs. When Validation AS receives the messages, it MUST conduct ROV on the messages and only stores the target source prefixes with the "valid" ROV state. The "Unknown" and "Invalid" target source prefixes will be ignored. In Figure 3, P1 MUST be authorized to AS 1, and then AS 1 advertises its own AS number and P1 to AS 4 through an SPA message.
 - Validation AS can also obtain the target source prefixes directly from RPKI ROAs or other RPKI data.
- * SPD process: After SPA process, Source AS can send SPD messages to Validation AS for notifying the wanted incoming directions of target source prefixes. That is, Source AS can specify from which neighbor ASes of Validation AS the target source prefixes will arrive. Validation AS will learn the specified incoming directions of target source prefixes and will use prefix blocklist for denying the target source prefixes coming from unwanted directions (neighbor ASes). The wanted incoming directions of target source prefixes can be obtained via the following methods for different purposes:
 - Automatically obtained from the RIB of Source AS. In Figure 3, AS 1 can specify that AS 2 and AS 3 are the wanted incoming directions of P1. AS 4 will block the packets with source addresses of P1 coming from neighbor ASes of AS 5 and AS 6. The use cases can be 1) proactive SAV for customer's prefixes or 2) key source address's forwarding path protection (i.e., keeping control plane path and data plane path consistent).
 - Obtained from security center of Source AS or Validation AS. Security center can detect source address-spoofed DDoS attacks and disseminates rules through BGP SAVNET to reactively filter source address at specific interfaces for mitigating DDoS suffered by customers.
 - Obtained from RPKI ASPA records or other RPKI data.

4. BGP SAVNET Peering Models

4.1. Full-mesh IBGP Peering

This peering model is for both intra- and inter-domain BGP SAVNET. In this model, 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 5.

4.2. EBGP Peering between ASes

Inter-domain BGP SAVNET requires eBGP sessions which can be single-hop or multi-hop. In this peering model, for the AS enabling BGP SAVNET, at least one border router in Source AS MUST establish the BGP SAVNET sessions with the border routers in Validation AS. SAVNET messages between ASes will be advertised through these sessions. The extensions of BGP messages for carrying SAVNET messages will be introduced in Section 5.

5. BGP SAVNET Protocol Extension

5.1. BGP SAVNET SAFI

To make good isolation with existing BGP services, this section defines BGP SAVNET SAFIs under the IPv4 address family and the IPv6 address family, respectively. The values require IANA registration as specified in Section 11. 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.

5.2. 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.

5.2.1. SPA TLVs for Intra-domain BGP SAVNET

The BGP SAVNET NLRI TLV each carries a 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:

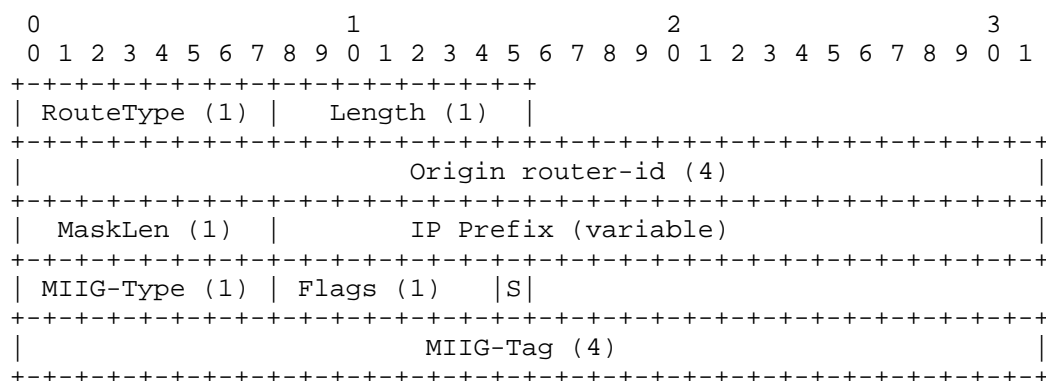


Figure 4: SPA TLV format

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 subnets. It can be a local prefix or a local domain prefix.

- Type value 1: Single-homing interface. Indicates that this prefix comes from a subnet that is single-homed to the local domain.
 - Type value 2: Complete multi-homing interface. Indicates that this prefix comes from a subnet that is multi-homed to the local domain, and is connected only to the local domain.
 - Type value 3: Incomplete multi-homing interface. Indicates that this prefix comes from a subnet that is multi-homed to the local domain and other domains.
 - Type value 4: Internet interface. Indicates that this prefix comes from a interface that is connected to the Internet.
 - Type value 5~255: Reserved for future use.
 - Notes: The type values of 3 and 4 are pre-defined for future use, and they should not appear in SPA TLVs (i.e., no need to advertise the prefixes from the interfaces of Type 3 and Type 4).
- * 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 owned by one subnet. The value of 0 indicates that the SPA prefix is not owned by only one subnet.
- * MIIG-Tag (non-key): Multi-homing Ingress Interface Group Tag. The value ranges from 1 to 0xFFFFFFFFE. The value 0 is invalid and the value 0xFFFFFFFF is reserved.

5.2.2. SPA TLVs for Inter-domain BGP SAVNET

This type of TLVs are used in SPA process between ASes.

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| RouteType (1) | Length (1) |
+-----+-----+-----+-----+
|                               Source AS Number (4)                               |
+-----+-----+-----+-----+
| MaskLen (1) | IP Prefix (variable) |
+-----+-----+-----+-----+
| Flags (1) |
+-----+-----+-----+-----+

```

Figure 5: SPA TLV format

The meaning of these fields are as follows:

- * RouteType (key): Type of the BGP SAVNET NLRI TLV, the value is 2 for SPA TLV between ASes.
- * Length: The length of the BGP SAVNET NLRI value, the RouteType and Length fields are excluded.
- * Source AS Number (key): The AS number of the originating AS of this advertised source prefix.
- * 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.
- * Flags (non-key): Reserved for future use.

5.3. BGP SAVNET Refresh

Two BGP SAVNET speakers MUST exchange Route Refresh Capability [RFC2918] to ensure that they are both capable of processing the SPD message carried in the BGP Refresh message.

The SPD TLV is carried in a BGP Refresh message after the BGP Refresh message body, as follows:

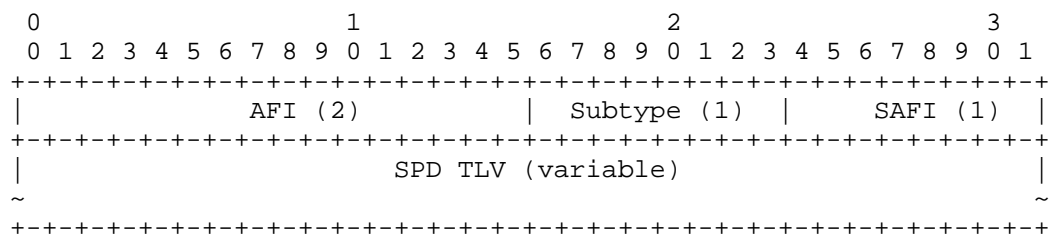


Figure 6: BGP-REFRESH with SPD TLV format

The AFI field is either 1 (IPv4) or 2 (IPv6). The SAFI field is the newly defined SAVNET SAFI. The Subtype field should be a new value assigned to SAVNET [RFC7313]. By carrying an SPD TLV, a BGP SAVNET Refresh message MUST NOT be processed as a Route-Refresh (as a re-advertisement request) and SHOULD only be used in the SPD process. A BGP SAVNET Refresh message without an SPD TLV SHOULD be processed as a Route-Refresh as defined in Route Refresh Capability [RFC2918].

5.3.1. The SPD TLVs for Inter-domain BGP SAVNET

This type of TLVs are used in SPD process between ASes.

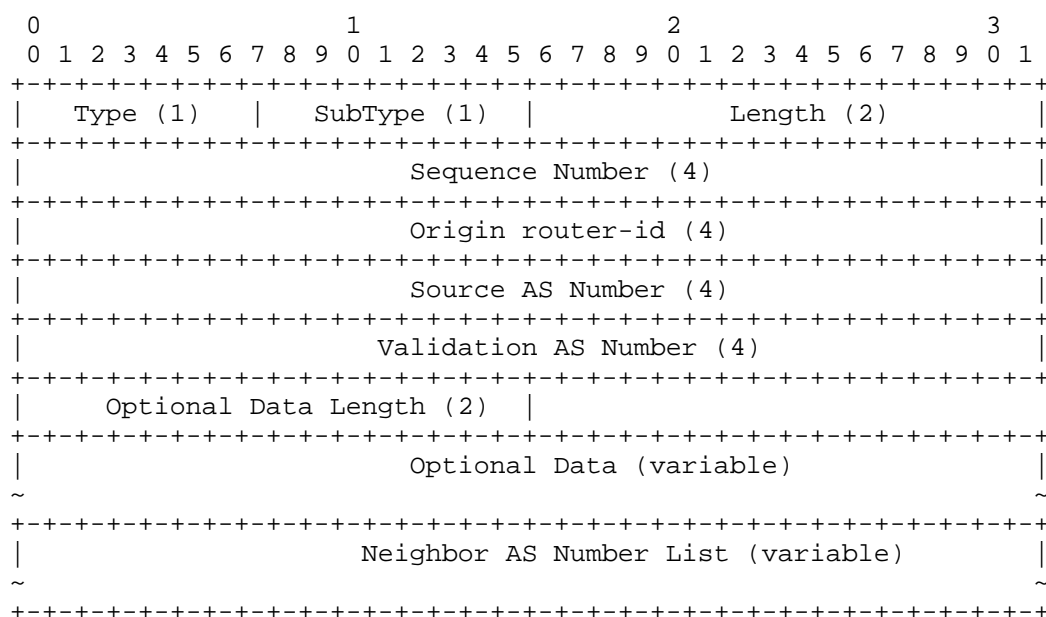


Figure 7: SPD TLV format

The meaning of these fields are as follows:

- * Type: TLV Type, the value is 2 for SPD TLV.
- * SubType: TLV Sub-Type, value is 2 for SPD TLV between an AS.
- * Length: The length of the SPD TLV value, the Type, SubType and Length fields are excluded.
- * Sequence Number: Indicates the sequence of Source Path Discovery process. The initial value is 0 and the value increases monotonically.

- * Origin router-id: The router ID of the originating node of the Source Path Discovery process.
- * Source AS Number (key): The AS number of the source AS whose source prefixes need to be validated in the validation AS.
- * Validation AS Number (key): The AS number of the validation AS who conducts validation for the source prefixes of source AS.
- * Optional Data Length: The length of the optional data field in bytes. The value can be 0 when there is no optional data.
- * Optional Data: Reserved for future use.
- * Neighbor AS Number List: List of neighbor AS, from which the validation AS will receive the data packets with the source prefixes of the source AS.

6. Decision Process with BGP SAVNET

The Decision Process described in [RFC4271] works to determine 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 and SPD 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.

6.1. 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.

6.1.1. 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.

7. Error Handling

7.1. 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.

7.2. 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.

7.3. Process of BGP SAVNET Refresh

Each BGP Refresh message MUST contain at most one SPD TLV. When a BGP SAVNET speaker receives a BGP Refresh packet with multiple SPD TLVs, only the first one SHOULD be processed.

7.4. Process of BGP SAVNET SPD TLVs

When a BGP SAVNET speaker receives an SPD TLV with an undefined type or subtype, it SHOULD be ignored.

When a BGP SAVNET speaker receives an SPD 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 SPD TLV with a validation AS number, 0 source AS number, AS_TRANS number (23456), or the source AS number equals the validation AS number, it MUST be considered malformed.

When a BGP SAVNET speaker receives an SPD TLV with an optional data sub-TLV that is an undefined type, it SHOULD be ignored.

When a BGP SAVNET speaker receives an SPD TLV with a DestList field that is not a multiple of 4 in length, it MUST be considered malformed.

When a BGP SAVNET speaker receives a Refresh message with a malformed SPD TLV, it MUST ignore the received message. When discarding a malformed message, a BGP SAVNET speaker MAY log a specific error.

When a BGP SAVNET speaker receives an SPD TLV with a sequence number that does not match the local recorded sequence number:

- * If the newly received sequence number is numerically larger, the local recorded sequence number SHOULD be updated to the newly received sequence number.
- * If the newly received sequence number is numerically smaller, the local recorded sequence number SHOULD NOT be updated, and the BGP SAVNET speaker SHOULD log a specific error.

8. Convergence Considerations

The convergence process of BGP SAVNET is relatively simple. First, the convergence process is mainly the message propagation process. BGP SAVNET messages should have similar propagation speed to normal routes. Second, BGP SAVNET supports independent and incremental update. Routers enable SAVNET can update local SAV rules immediately and there is no need to wait for complete information updates.

9. Deployment Considerations

Both intra- and inter-domain BGP SAVNET have good deployability. For intra-domain BGP SAVNET, upgrading part of routers can also work well. For example, only upgrade the routers (two or more) multi-homed by the same subnet, or upgrade one edge router and one border router. With more routers getting deployed, the network can get more protection. For inter-domain BGP SAVNET, any pair of ASes can upgrade and work well. There is no dependence on other ASes.

10. Security Considerations

Security problems are mainly in inter-domain scenarios.

- * For communication security, inter-domain BGP SAVNET takes a point-to-point communication model and thus has a simple trust model. The communication between source AS and validation AS can be protected by TLS.
- * For content security, the advertised source prefixes of Source AS MUST be authorized to Source AS by RPKI ROAs. When Validation AS receives the messages, it MUST conduct ROV on the messages.

11. IANA Considerations

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

Acknowledgements

Many thanks to the contributions from Fang Gao. Also thanks for the comments from Jeff Haas, Antoin Verschuren, Zhibin Dai, Keyur Patel, Shunwan Zhuang, David Lamparter, etc.

References

Normative References

- [RFC2918] Chen, E., "Route Refresh Capability for BGP-4", RFC 2918, DOI 10.17487/RFC2918, September 2000, <<https://www.rfc-editor.org/info/rfc2918>>.
- [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>>.
- [I-D.li-savnet-intra-domain-architecture]
Li, D., Wu, J., Qin, L., Geng, N., Chen, L., Huang, M., and F. Gao, "Intra-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-li-savnet-intra-domain-architecture-07, 16 March 2024, <<https://datatracker.ietf.org/doc/html/draft-li-savnet-intra-domain-architecture-07>>.
- [I-D.wu-savnet-inter-domain-architecture]
Li, D., Chen, L., Geng, N., Liu, L., and L. Qin, "Inter-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-wu-savnet-inter-domain-architecture-12, 3 November 2024, <<https://datatracker.ietf.org/doc/html/draft-wu-savnet-inter-domain-architecture-12>>.
- [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>>.
- [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>>.
- [RFC7313] Patel, K., Chen, E., and B. Venkatachalapathy, "Enhanced Route Refresh Capability for BGP-4", RFC 7313, DOI 10.17487/RFC7313, July 2014, <<https://www.rfc-editor.org/info/rfc7313>>.
- [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-22, 19 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-22>>.
- [I-D.ietf-savnet-inter-domain-problem-statement]
Li, D., Qin, L., Liu, L., Huang, M., and K. Sriram, "Gap Analysis, Problem Statement, and Requirements for Inter-Domain SAV", Work in Progress, Internet-Draft, draft-ietf-savnet-inter-domain-problem-statement-15, 18 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-problem-statement-15>>.

[I-D.huang-savnet-sav-table]

Huang, M., Cheng, W., Li, D., Geng, N., Liu, Chen, L., and
C. Lin, "General Source Address Validation Capabilities",
Work in Progress, Internet-Draft, draft-huang-savnet-sav-
table-08, 10 December 2024,
<<https://datatracker.ietf.org/doc/html/draft-huang-savnet-sav-table-08>>.

[manrs-antispoofing]

"MANRS Implementation Guide", January 2023,
<<https://www.manrs.org/netops/guide/antispoofing>>.

Authors' Addresses

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

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

Zhen Tan
Huawei Technologies
Beijing
China
Email: tanzhen6@huawei.com

Mingxing Liu
Huawei Technologies
Beijing
China
Email: liumingxing7@huawei.com

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