

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: 5 March 2026

D. Li
Tsinghua University
L. Chen
Zhongguancun Laboratory
N. Geng
Huawei
L. Liu
L. Qin
Zhongguancun Laboratory
1 September 2025

Inter-domain Source Address Validation (SAVNET) Architecture
draft-ietf-savnet-inter-domain-architecture-02

Abstract

This document introduces an inter-domain SAVNET architecture for performing AS-level SAV and provides a comprehensive framework for guiding the design of inter-domain SAV mechanisms. The proposed architecture empowers ASes to generate SAV rules by sharing SAV-specific information between themselves, which can be used to generate more accurate and trustworthy SAV rules in a timely manner compared to the general information. During the incremental or partial deployment of SAV-specific information, it can utilize general information to generate SAV rules, if an AS's SAV-specific information is unavailable. Rather than delving into protocol extensions or implementations, this document primarily concentrates on proposing SAV-specific and general information and guiding how to utilize them to generate SAV rules. To this end, it also defines some architectural components and their relations.

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 5 March 2026.

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	3
1.1. Requirements Language	5
2. Terminology	5
3. Design Goals	7
4. Inter-domain SAVNET Architecture Overview	7
5. SAV-related Information	12
5.1. General Information	12
5.1.1. RPKI ROA objects and ASPA Objects	12
5.1.2. Local Routing Information	13
5.1.3. IRR Data	13
5.2. SAV-specific Information	13
6. SAV Information Base	14
7. SAVNET Communication Mechanism	17
7.1. SAV-specific Information Communication Mechanism	18
7.2. General Information Communication Mechanism	21
7.3. Management Information Communication Mechanism	21
8. Use Cases	22
8.1. SAV at Customer Interfaces	22
8.1.1. Limited Propagation of Prefixes	22
8.1.2. Hidden Prefixes	24
8.1.3. Reflection Attacks	25
8.1.4. Direct Attacks	26
8.2. SAV at Provider/Peer Interfaces	27
8.2.1. Reflection Attacks	27
8.2.2. Direct Attacks	29
9. Meeting the Design Requirements of Inter-domain SAVNET	30
9.1. Improving Validation Accuracy over Existing Mechanisms	30
9.2. Working in Incremental/Partial Deployment	30
9.3. Reducing Operational Overhead	32
9.4. Guaranteeing Convergence	32
9.5. Providing Necessary Security Guarantee	33
10. Manageability Considerations	35

11. Privacy Considerations	35
12. IANA Considerations	35
13. Scope and Assumptions	35
14. Contributors	37
15. References	37
15.1. Normative References	37
15.2. Informative References	38
Acknowledgements	39
Authors' Addresses	39

1. Introduction

Attacks based on source IP address spoofing, such as reflective DDoS and flooding attacks, continue to present significant challenges to Internet security. Mitigating these attacks in inter-domain networks requires effective source address validation (SAV). While BCP84 [RFC3704] [RFC8704] offers some SAV solutions, such as ACL-based ingress filtering and uRPF-based mechanisms, existing inter-domain SAV mechanisms have limitations in terms of validation accuracy and operational overhead in different scenarios [inter-domain-ps].

There are various existing general information from different sources including RPKI ROA objects and ASPA objects, RIB, FIB, and Internet Routing Registry (IRR) data, which can be used for inter-domain SAV. Generating SAV rules based on general information, however, cannot well satisfy the requirements for new inter-domain SAV mechanisms proposed in [inter-domain-ps]. As analyzed in Section 5, general information from RPKI ROA objects and ASPA objects can be used to infer the prefixes and their permissible incoming directions yet cannot be updated in a timely manner to adapt to the prefix or route changes, and the local routing information, which represents the general information from RIB or FIB, cannot deal with the asymmetric routing scenarios and may lead to improper blocks or improper permits, while IRR data do not update in a timely manner either and are not always accurate.

Consequently, to address these issues, the inter-domain SAVNET architecture focuses on providing a comprehensive framework and guidelines for the design and implementation of new inter-domain SAV mechanisms. Inter-domain SAVNET architecture proposes SAV-specific information and uses it to generate SAV rules. SAV-specific information consists of prefixes and their corresponding legitimate incoming direction to enter an AS. Inter-domain SAVNET architecture can use it to generate more accurate SAV rules. In order to gather the SAV-specific information, a SAV-specific information communication mechanism would be developed for origination, processing, propagation, and termination of the messages which carry the SAV-specific information, and it can be implemented by a new

protocol or extending an existing protocol. When the prefixes or routes change, it can update the SAV-specific information automatically in a timely manner. Also, the inter-domain SAVNET architecture will communicate the SAV-specific information over a secure connection between authenticated ASes.

Moreover, during the incremental/partial deployment period of the SAV-specific information, the inter-domain SAVNET architecture can leverage the general information to generate SAV rules, if the SAV-specific information of an AS is unavailable. Multiple information sources may exist concurrently, to determine the one used for generating SAV rules, the inter-domain SAVNET architecture assigns priorities to the SAV-specific information and different general information and generates SAV rules using the SAV-related information with the highest-priority. SAV-specific information has the highest priority and the priorities of RPKI ROA objects and ASPA objects, RIB, FIB, and IRR data decrease in turn.

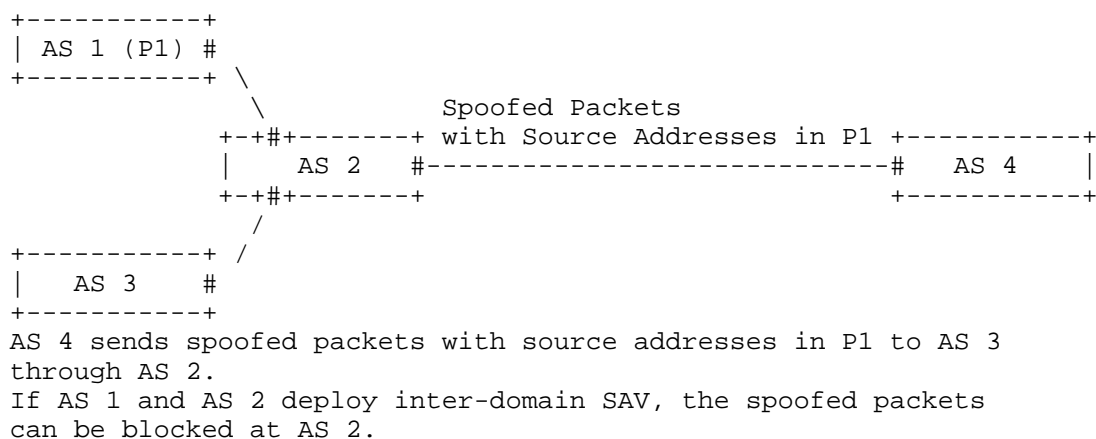


Figure 1: An example for illustrating the incentive of deploying inter-domain SAVNET architecture.

The inter-domain SAVNET architecture provides the incentive to deploy inter-domain SAV for operators. Figure 1 illustrates this using an example. P1 is the source prefix of AS 1, and AS 4 sends spoofing packets with P1 as source addresses to AS 3 through AS 2. Assume AS 4 does not deploy intra-domain SAV, these spoofing packets cannot be blocked by AS 4. Although AS 1 can deploy intra-domain SAV to block incoming packets which spoof the addresses of AS 1, these spoofing traffic from AS 4 to AS 3 do not go through AS 1, so they cannot be blocked by AS 1. Inter-domain SAVNET architecture can help in this scenario. If AS 1 and AS 2 deploy inter-domain SAVNET architecture, AS 2 knows that the packets with P1 as source addresses should come

from AS 1, and the spoofing packets can thus be blocked by AS 2 since they come from the incorrect direction. Specifically, by proposing SAV-specific information and using it to generate SAV rules, the inter-domain SAVNET architecture gives more deployment incentive compared to existing inter-domain SAV mechanisms, which will be analyzed in Section 8.

In addition, this document primarily proposes a high-level architecture for describing the communication flow of SAV-specific information and general information, guiding how to utilize the SAV-specific information and general information for generating SAV rules and deploy an inter-domain SAV mechanism between ASes. This document does not specify protocol extensions or implementations. Its purpose is to provide a conceptual framework and guidance for the design and development of inter-domain SAV mechanisms, allowing implementers to adapt and implement the architecture based on their specific requirements and network environments.

1.1. 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. Terminology

SAV Rule:

The rule that indicates the validity of a specific source IP address or source IP prefix.

SAV Table:

The table or data structure that implements the SAV rules and is used for performing source address validation on the data plane.

SAV-specific Information:

The information that is specialized for SAV rule generation, includes the source prefixes and their legitimate incoming directions to enter an AS, and is gathered by the communication between ASes with the SAV-specific information communication mechanism.

SAV-specific Information Communication Mechanism:

The mechanism that is used to communicate SAV-specific information between ASes and can be implemented by a new protocol or an extension to an existing protocol.

Local Routing Information:

The information that is stored in ASBR's local RIB or FIB and can be used to generate SAV rules in addition to the routing purpose.

General Information:

The information that is not specialized for SAV but can be utilized to generate SAV rules, and is initially utilized for other purposes. Currently, the general information consists of the information from RPKI ROA objects and ASPA objects, local routing information, and the one from IRR data.

SAV-related Information:

The information that can be used to generate SAV rules and includes SAV-specific information and general information.

SAVNET Agent:

The agent within a SAVNET-adopting AS that is responsible for gathering SAV-related information and utilizing it to generate SAV rules.

SAV Information Base:

SAV information base is a table or data structure for storing SAV-related information collected from different SAV information sources and is a component within SAVNET agent.

SAV Information Base Manager:

SAV information base manager maintains the SAV-related information in the SAV information base and uses it to generate SAV rule accordingly, and is a component within SAVNET agent.

Improper Block:

The validation results that the packets with legitimate source addresses are blocked improperly due to inaccurate SAV rules.

Improper Permit:

The validation results that the packets with spoofed source addresses are permitted improperly due to inaccurate SAV rules.

Source AS:

The AS which deploys SAVNET agent and communicates its own SAV-specific information to other ASes for generating SAV rules.

Validation AS:

The AS which deploys SAVNET agent and generates SAV rules according to the received SAV-specific information from source ASes.

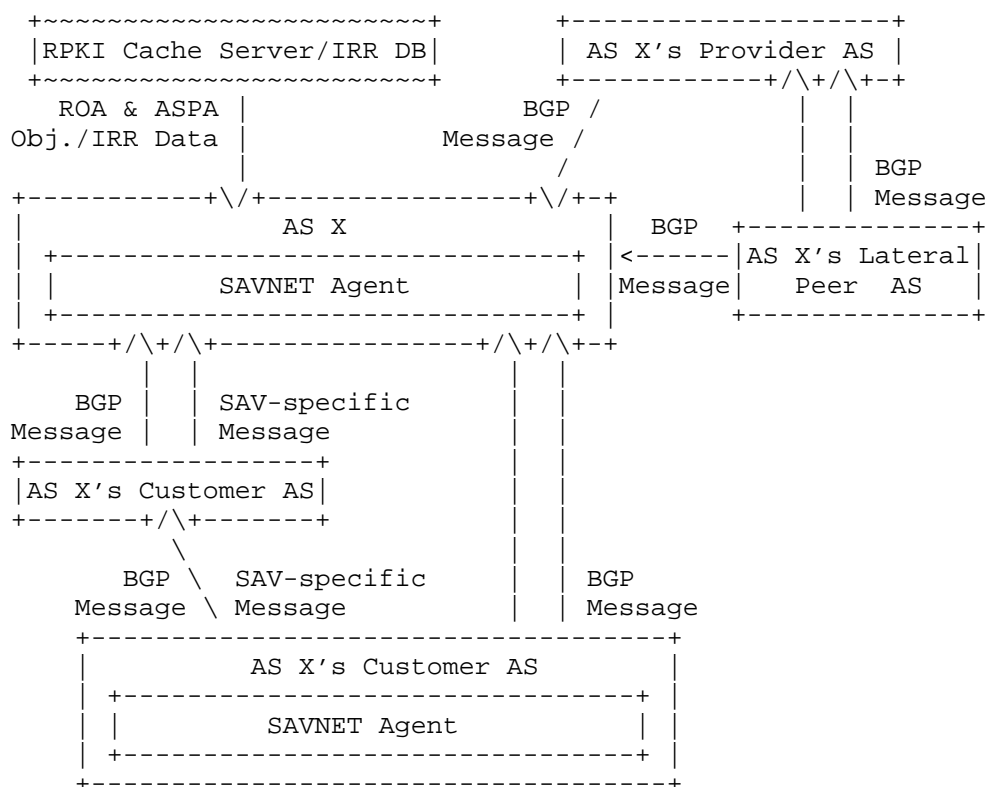
3. Design Goals

The inter-domain SAVNET architecture aims to improve SAV accuracy and facilitate partial deployment with low operational overhead, while guaranteeing convergence and providing security guarantees to the communicated information, which corresponds to the requirements for new inter-domain SAV mechanisms proposed in the inter-domain SAVNET architecture draft [inter-domain-ps]. The overall goal can be broken down into the following aspects:

- * *G1*: The inter-domain SAVNET architecture should learn the real paths of source prefixes to any destination prefixes or permissible paths that can cover their real paths, and generate accurate SAV rules automatically based on the learned information to avoid improper blocks and reduce improper permits as much as possible.
- * *G2*: The inter-domain SAVNET architecture should provide sufficient protection for the source prefixes of ASes that deploy it, even if only a portion of the Internet does the deployment.
- * *G3*: The inter-domain SAVNET architecture should adapt to dynamic networks and asymmetric routing scenarios automatically.
- * *G4*: The inter-domain SAVNET architecture should promptly detect the network changes and launch the convergence process in a timely manner, while reducing improper blocks and improper permits during the convergence process.
- * *G5*: The inter-domain SAVNET architecture should provide security guarantees for the communicated SAV-specific information.

Other design goals, such as low operational overhead and easy implementation, are also very important and should be considered in specific protocols or protocol extensions.

4. Inter-domain SAVNET Architecture Overview



AS X and one of its customer ASes have deployed SAVNET agent and can exchange SAV-specific information with each other.

Figure 2: Inter-domain SAVNET architecture.

Figure 2 provides an overview of the inter-domain SAVNET architecture, showcasing an AS topology and the flow of SAV-related information among ASes. The topology captures the full spectrum of AS relationships in the Internet, displaying all peer ASes of AS X including customers, lateral peers, and providers and the existence of multiple physical links between ASes. Arrows in the figure indicate the direction of the corresponding SAV-related information from its source to AS X, such as gathering RPKI ROA objects and ASPA objects from RPKI cache server. The inter-domain SAVNET architecture conveys the SAV-related information through various mediums such as SAV-specific messages, BGP messages, RTR messages, and FTP messages. Based on the SAV-related information, AS X generates SAV rules. It is also worth noting that the inter-domain SAVNET architecture discusses AS-level inter-domain SAV.

Figure 2 uses AS X as the representative to illustrate that what SAV-related information the SAVNET agent within AS X will collect and where the information is from. AS X has deployed SAVNET agent and can generate SAV rules to perform inter-domain SAV by consolidating the SAV-related information. It can obtain SAV-specific information from its customer AS which deploys SAVNET agent and local routing information originating from the BGP update messages of its neighbor ASes. Also, AS X can obtain RPKI ROA objects and ASPA objects from RPKI cache server and IRR data from IRR database.

The inter-domain SAVNET architecture proposes SAV-specific information, which is more accurate and trustworthy than existing general information, and can update in a timely manner. SAV-specific information consists of prefixes and their legitimate incoming directions. The SAVNET agent communicates SAV-specific information between ASes via SAV-specific messages, when prefixes or routes change, it can launch SAV-specific messages timely to update SAV-specific information. Additionally, when SAVNET agent receives SAV-specific messages, it will validate whether the SAV-specific connections for communicating SAV-specific messages are authentic connections from authenticated ASes. Therefore, when SAV-specific information of an AS is available, SAVNET agent will use it to generate SAV rules.

Furthermore, if the SAV-specific information is needed to communicate between ASes, a new SAV-specific information communication mechanism would be developed to exchange the SAV-specific messages between ASes which carry the SAV-specific information. It should define the data structure or format for communicating the SAV-specific information and the operations and timing for originating, processing, propagating, and terminating the SAV-specific messages.

The SAVNET agent should launch SAV-specific messages to adapt to the route changes in a timely manner. The SAV-specific information communication mechanism should handle route changes carefully to avoid improper blocks. The reasons for leading to improper blocks may include late detection of route changes, delayed message transmission, or packet losses. During the convergence process of the SAV-specific information communication mechanism, the inter-domain SAVNET architecture can use the information from RPKI ROA objects and ASPA objects to generate SAV rules until the convergence process is finished, since these information includes topological information and is more stable, and can thus avoid improper blocks. However, the detailed design of the SAV-specific information communication mechanism for dealing with route changes is outside the scope of this document.

In the incremental/partial deployment stage of the inter-domain SAVNET architecture, when the SAV-specific information of some ASes is unavailable, SAVNET agent can leverage general information to generate SAV rules. If all these general information is available, it is recommended to use RPKI ROA objects and ASPA objects to generate SAV rules. Since compared to the local routing information and IRR data, they can provide authoritative prefixes and topological information and have less improper blocks. The systematic recommendations for the utilizations of SAV-related information and the corresponding rationale will be illustrated in Section 6.

SAV-specific information communication mechanism will require specifying a new inter-router (or inter-AS) communication protocol or modifying an existing one. Therefore, while this is pursued, a new SAV mechanism that utilizes RPKI objects (e.g., ROA, ASPA) and BGP data (RIB/FIB) can be pursued. The latter solution may have the potential for deployment in the near term since it utilizes existing SAV-related information and can be deployed by network operators by policy configuration on routers.

For ASes that support a network controller, such as a multi-AS or single-AS controller, operators can deploy the SAVNET agent on the controller to represent the ASes it manages and communicate SAV-specific information with others. Additionally, ASes managed by the same controller may obtain SAV-specific information directly from the controller without needing to communicate with each other.

Regarding the security concerns, the inter-domain SAVNET architecture shares the similar security threats with BGP and can leverage existing BGP security mechanisms to enhance both session and content security.

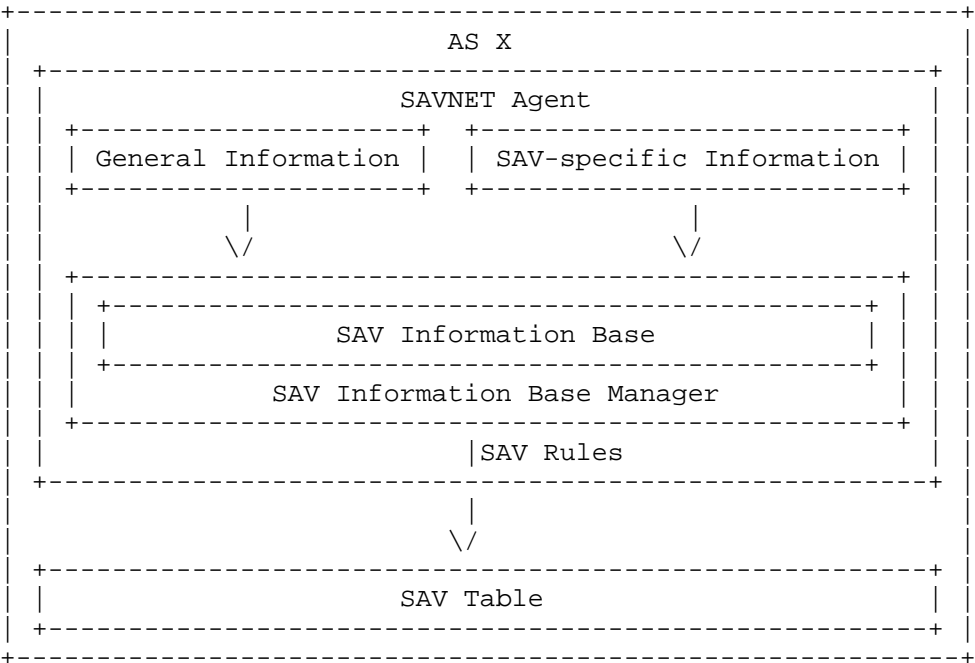


Figure 3: SAVNET agent and SAV table within AS X in Figure 2.

Figure 3 displays the SAVNET agent and SAV table within AS X. The SAVNET agent can obtain the SAV-specific information and general information from various SAV information sources including SAV-specific messages from other ASes, RPKI cache server, and RIB or FIB as long as they are available. The SAV information base (SIB) within the SAVNET agent can store the SAV-specific information and general information and is maintained by the SIB manager. And SIB manager generates SAV rules based on the SIB and fills out the SAV table on the data plane. Moreover, the SIB can be managed by network operators using various methods such as YANG [RFC6020], Command-Line Interface (CLI), remote triggered black hole (RTBH) [RFC5635], and Flowspec [RFC8955]. The detailed collection methods of the SAV-related information depend on the deployment and implementation of the inter-domain SAV mechanisms and are out of scope for this document.

In the data plane, the packets coming from other ASes will be validated by the SAV table and only the packets which are permitted by the SAV table will be forwarded to the next hop. To achieve this, the router looks up each packet's source address in its local SAV table and gets one of three validity states: "Valid", "Invalid" or "Unknown". "Valid" means that there is a source prefix in SAV table

covering the source address of the packet and the valid incoming interfaces covering the actual incoming interface of the packet. According to the SAV principle, "Valid" packets will be forwarded. "Invalid" means there is a source prefix in SAV table covering the source address, but the incoming interface of the packet does not match any valid incoming interface so that such packets will be dropped. "Unknown" means there is no source prefix in SAV table covering the source address. The packet with "unknown" addresses can be dropped or permitted, which depends on the choice of operators. The structure and detailed usage of SAV table can refer to [sav-table].

5. SAV-related Information

SAV-related information represents the information that can be used for SAV and consists of RPKI ROA objects and ASPA objects, local routing information, IRR data, and SAV-specific information. In the inter-domain SAVNET architecture, RPKI ROA objects and ASPA objects, local routing information, and IRR data are categorized into general information. In the future, if a new information source is created and can be used for SAV, but is not originally and specially used for SAV, its information can be categorized into general information. In other words, general information can also be considered as dual-use information.

5.1. General Information

General information refers to the information that is not directly designed for SAV but can be utilized to generate SAV rules, and includes RPKI ROA objects and ASPA objects, local routing information, and IRR data.

5.1.1. RPKI ROA objects and ASPA Objects

The RPKI ROA objects and ASPA objects are originally designed for the routing security purpose. RPKI ROA objects consists of {prefix, maximum length, origin AS} information and are originally used to mitigate the route origin hijacking, while RPKI ASPA objects consists of {ASN, Provider AS Set} information and are originally used to mitigate the route leaks. Both the objects are verified and authoritative. They are also stable and will not be updated frequently.

Based on ASPA objects, the AS-level network topology can be constructed. And according to the ROA objects and the constructed AS-level topology information, an AS can learn all the permissible paths of the prefixes from its customer cone. Therefore, the prefixes and all its permissible incoming directions can be obtained.

SAV based on RPKI ROA and ASPA objects only may lead to improper blocks, because not all ASes register their ROA and ASPA objects at the current stage. In addition, all the permissible incoming directions learned from ASPA objects do not only consist of the real incoming directions of the prefixes, but also the extra non-used incoming directions by the legitimate traffic. Only utilizing RPKI ROA and ASPA objects for generating SAV rules with their full deployment within the customer cone may lead to improper permits when not all ASes perform SAV.

According to a recent study [rpki-time-of-flight], the process of updating RPKI information typically requires several minutes to an hour. This encompasses the addition or deletion of RPKI objects and the subsequent retrieval of updated information by ASes.

5.1.2. Local Routing Information

The local routing information is originally used to guide the packet forwarding on each router and can be stored in the local RIB or FIB. It can be parsed from the BGP messages communicated between ASes. Existing uRPF-based SAV mechanisms [RFC3704] [RFC8704] use the local routing information to generate SAV rules. As analyzed in [inter-domain-ps], in the asymmetric routing scenarios, these mechanisms, generating SAV rules using local routing information only, have accuracy problems and would lead to improper permits or improper blocks.

5.1.3. IRR Data

The IRR data consist of ASes and their corresponding prefixes and can be used for SAV [RFC8704]. However, SAV using IRR data only would have limited functioning scope, in inter-domain networks, it may only be able to prevent spoofing by a stub AS. In addition, the IRR data are not always accurate [RFC8704].

5.2. SAV-specific Information

SAV-specific information is the information that is specifically designed for SAV and consists of prefixes and their legitimate incoming directions to enter ASes. It can be contained in the SAV-specific messages which are communicated between ASes which deploy the inter-domain SAVNET architecture. When parsing the SAV-specific messages and obtaining the SAV-specific information, ASes can learn the prefixes and their legitimate incoming direction to enter themselves.

Moreover, in the inter-domain SAVNET architecture, a SAV-specific information communication mechanism is used to communicate SAV-specific information between ASes and distribute the updated information to the relative ASes automatically in a timely manner once the prefixes or routes change. Compared against general information, it may be expected that SAV-specific information is more accurate and trustworthy, while it can update the SAV rules in a timely manner to adapt to the prefix or route changes.

6. SAV Information Base

SAV Information Sources		Trustworthiness	SAV Usage
SAV-specific Information		Specific security mechanism	
General Information	RPKI ROA & ASPA	X.509 certificates	Complementing each other.
	Local Routing Information	ROV, Route leak detection	
	IRR Data	Lower level of trust	

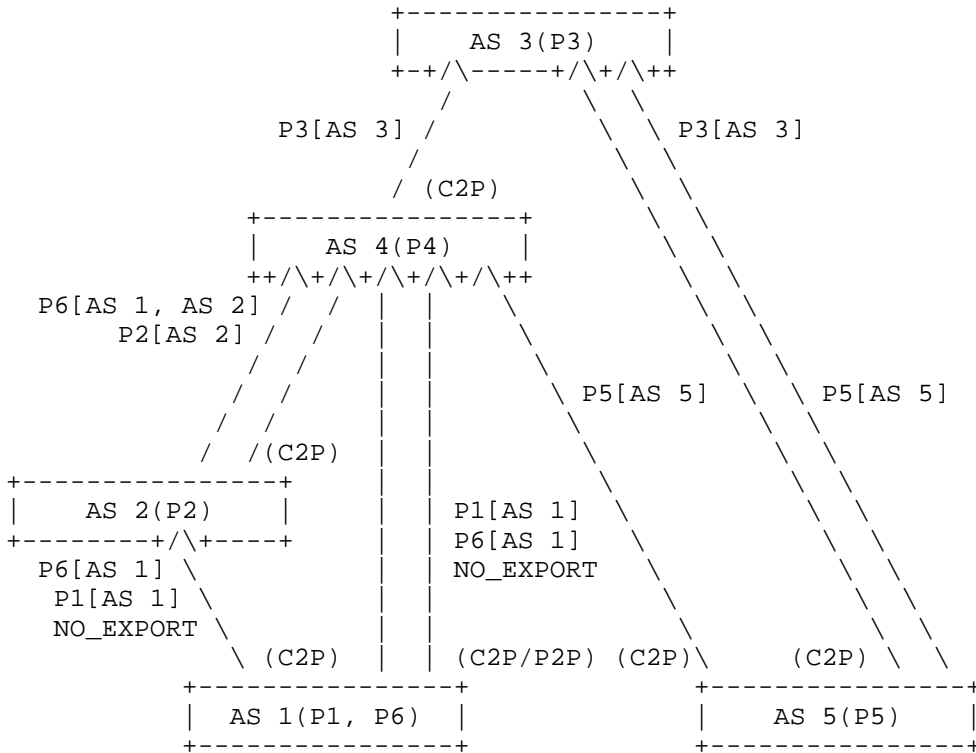
Figure 4: SAV information sources, their trustworthiness, and their usage for SAV.

The SIB is managed by the SIB manager, which can consolidate SAV-related information from different sources. Figure 4 presents the summary of SAV information sources, their trustworthiness, and the recommendation about how to use them for generating SAV rules. Inter-domain SAVNET architecture can use SAV-related information from different sources to generate SAV rules. It recommends using the information from all the available sources which is verified as trustworthy to generate SAV rules. Therefore, SAV information sources complement each other for SAV. For example, when both SAV-specific and general information are available, the inter-domain SAVNET would use them all equally.

The recommendation for completing each information source for SAV depends on their verified trustworthiness. Each AS deploying inter-domain SAVNET can utilize a specific security mechanism as discussed in Section 9.5 to validate the received SAV-specific information from other ASes to guarantee its Trustworthiness. RPKI ROA and ASPA objects use x.509 certificates [RFC5280] to protect the objects and guarantee the trustworthiness. Beside, local routing information relies on some techniques, such as ROV and route leak detection, e.g., ASPA or OTC, for achieving the Trustworthiness of their

information. Instead, IRR data has lower level of trust compared to others and may be disregarded if RPKI is deployed well, and they are usually updated in a slower manner than the real network changes and not always correct.

It is noteworthy that using each type of SAV information source complementally in Figure 4 is recommended rather than mandated. If a new inter-domain SAV mechanism needs to generate SAV rules using an information source, it should ensure that the correct information is obtained from the corresponding source and adopts appropriate SAV actions in the data plane to avoid improper block and minimize improper permit. A new inter-domain SAVNET mechanism, in line with the inter-domain SAVNET architecture, has the flexibility to determine the utilized SAV information sources and how to use them. Especially, when using RPKI ROA objects and ASPA objects as the SAV information source, the new inter-domain SAVNET mechanism should avoid jeopardizing the use of RPKI in routing security.



Both AS 1 and AS 4 deploy the inter-domain SAVNET architecture and can exchange the SAV-specific information with each other, while other ASes do not deploy it.

Figure 5: An example of AS topology.

Index	Prefix	Incoming Direction	Relation	SAV Information Source
0	P1	AS 2	Customer	SAV-specific Information
1	P1	AS 1	Customer	General Information
2	P2	AS 2	Customer	General Information
3	P3	AS 3	Provider	General Information
4	P5	AS 3	Provider	General Information
5	P5	AS 5	Customer	General Information
6	P6	AS 2	Customer	General Information SAV-specific Information
7	P6	AS 1	Customer	General Information

Figure 6: An example for the SAV information base of AS 4 in Figure 6.

We use the examples shown in Figure 5 and Figure 6 to introduce SIB and illustrate how to generate SAV rules based on the SIB. Figure 6 depicts an example of the SIB established in AS 4 displayed in Figure 5. Each row of the SIB contains an index, prefix, incoming direction of the prefix, relation between ASes, and the corresponding sources of the information. The incoming direction consists of customer, provider, and peer. For example, in Figure 6, the row with index 0 indicates the incoming direction of P1 is AS 2 and the information source is SAV-specific information. Note that the same SAV-related information may have multiple sources and the SIB records them all, such as the row indexed 6. Moreover, SIB should be carefully implemented in the specific protocol or protocol extensions to avoid becoming a heavy burden of the router, and the similar optimization approaches used for the RIB may be applied.

Recall that inter-domain SAVNET architecture generates SAV rules based on the SAV information sources in the SIB and each type of source complements each other. In addition, in the case of an AS's interfaces facing provider or lateral peer ASes where loose SAV rules are applicable, the inter-domain SAVNET architecture recommends to use blocklist at such directions to only block the prefixes that are sure not to come at these directions, while in the case of an AS's

interfaces facing customer ASes that necessitate stricter SAV rules, the inter-domain SAVNET architecture recommends to use allowlist to only permit the prefixes that are allowed to come at these directions.

Based on the above rules, taking the SIB in Figure 6 as an example to illustrate how the inter-domain SAVNET generates rules, AS 4 can conduct SAV as follows: SAV at the interfaces facing AS 3 blocks P1, P2, and P6 according to the rows indexed 0, 1, 2, 6, and 7 in the SIB, SAV at the interfaces facing AS 2 permits P1, P2, and P6 according to the rows indexed 0, 2, and 6 in the SIB, SAV at the interfaces facing AS 1 permit P1 and P6 according to the row indexed 1 and 7 in the SIB, and SAV at the interfaces facing AS 5 permits P5 according to the row indexed 5 in the SIB.

7. SAVNET Communication Mechanism

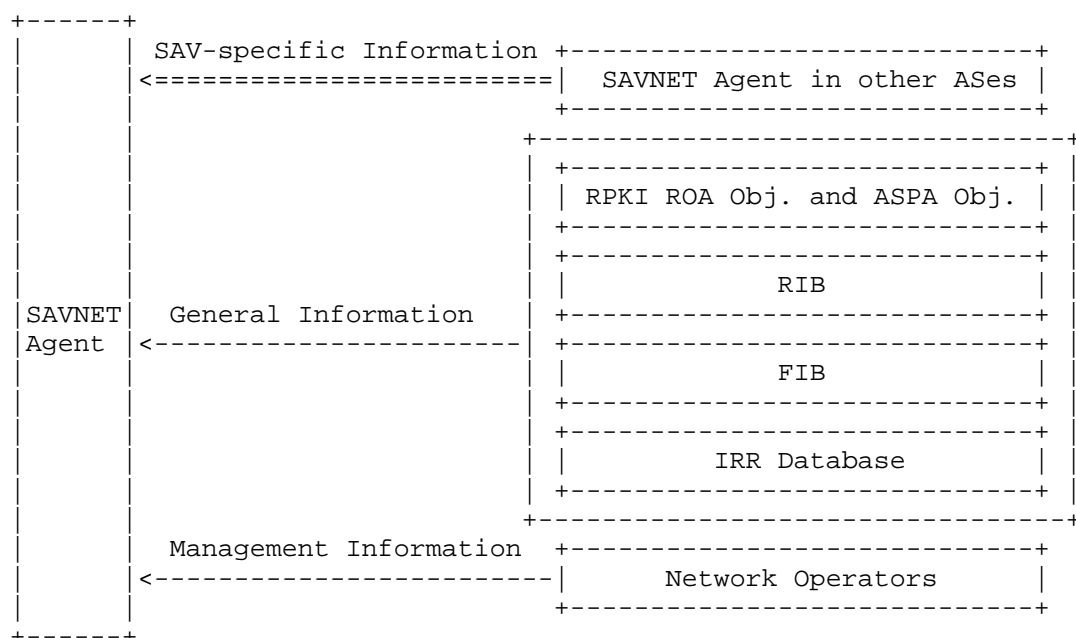
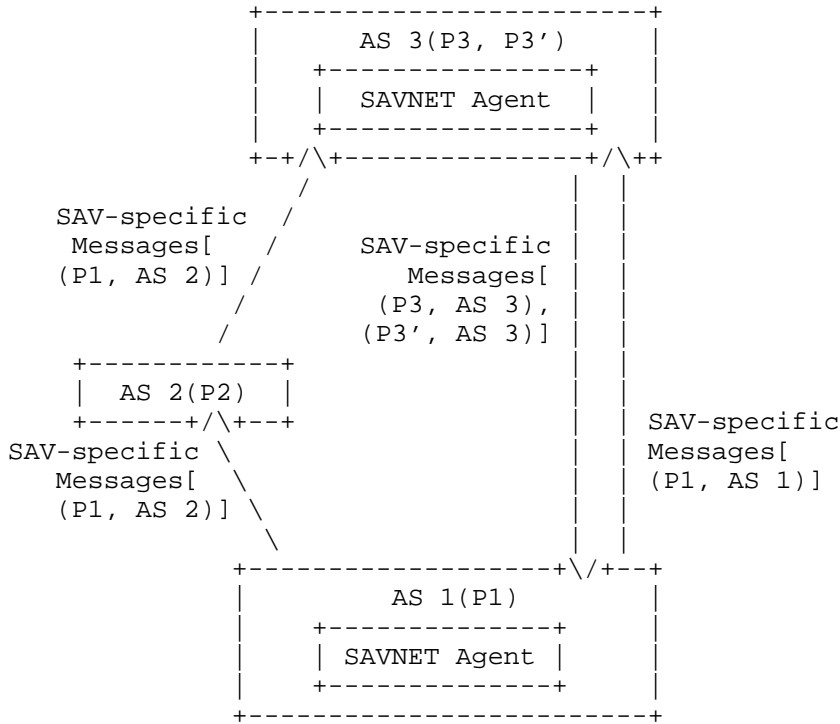


Figure 7: Gathering SAV-related information from different SAV information sources.

SAV-specific information relies on the communication between SAVNET agents, and general information can be from RPKI ROA objects and ASPA objects, RIB, FIB, and IRR data. Therefore, as illustrated in Figure 7, the SAVNET agent needs to receive the SAV-related information from these SAV information sources. SAVNET agent also

needs to accept the configurations from network operators for the management operations. Gathering these types of information relies on the SAVNET communication mechanism, which includes SAV-specific information communication mechanism, general information communication mechanism, and management information communication mechanism.

7.1. SAV-specific Information Communication Mechanism



- (1) The path of the legitimate traffic with source addresses in P1 and destination addresses in P3 is [AS 1, AS 2, AS 3].
- (2) The path of the legitimate traffic with source addresses in P1 and destination addresses in P3' is [AS 1, AS 3].
- (3) The path of the legitimate traffic with source addresses in P3 or P3' and destination addresses in P1 is [AS 3, AS 1].

Figure 8: An example for exchanging SAV-specific information with SAV-specific information communication mechanism between AS 1 and AS 3.

Figure 8 uses an example to show the exchange of SAV-specific information between AS 1 and AS 3 through SAV-specific messages. The SAV-specific information is expressed as <Prefix, Incoming Direction>

pairs, such as (P1, AS 1), (P1, AS 2), (P3, AS 3), and (P3', AS 3) in Figure 8. AS 1 needs to determine the incoming direction to AS 3 for its prefix P1 to obtain SAV-specific information, such as (P1, AS 1) and (P1, AS 2). It then assembles this information into SAV-specific messages to send to AS 3, which can subsequently generate SAV rules based on the received information. This process may require AS 1 to collaborate with intermediate ASes between AS 1 and AS 3 to obtain the SAV-specific information. Similarly, AS 3 needs to determine the incoming direction to AS 1 for its prefixes P3 and P3' and send the SAV-specific information, such as (P3, AS 3) and (P3', AS 3), to AS 1 through SAV-specific messages, allowing AS 1 to generate the corresponding SAV rules. AS 3 may also need to collaborate with intermediate ASes between AS 3 and AS 1 to gather the required information for obtaining the SAV-specific information.

The SAV-specific information can be exchanged between ASes via SAV-specific messages. SAV-specific messages are used to propagate or originate the SAV-specific information between ASes by the SAVNET agent. For an AS which initiates its own SAV-specific messages, its SAVNET agent needs to obtain the incoming direction of its own prefixes to enter other ASes and assemble them into the SAV-specific messages to the corresponding ASes. When ASes receive the SAV-specific messages, they parse the messages to obtain source prefixes and their corresponding incoming directions.

Additionally, if SAV-specific information is communicated between ASes, a new SAV-specific information communication mechanism would need to be developed to communicate it and can be implemented by a new protocol or extending an existing protocol. The SAV-specific information communication mechanism needs to define the data structure or format to communicate the SAV-specific messages and the operations and timing for originating, processing, propagating, and terminating the messages. If an extension to an existing protocol is used to exchange SAV-specific information, the corresponding existing protocol should not be affected. The SAVNET agent is the entity to support the SAV-specific communication mechanism. By parsing the SAV-specific messages, it obtains the prefixes and their incoming AS direction for maintaining the SIB. It is important to note that the SAVNET agent within an AS has the capability to establish connections with multiple SAVNET agents within different ASes, relying on either manual configurations by operators or an automatic mechanism. In addition, SAVNET agents should validate the authenticity of the connection for communicating the SAV-specific information to verify whether the SAV-specific information is provided over a secure connection with an authenticated AS.

The need for a SAV-specific communication mechanism arises from the facts that the SAV-specific information needs to be obtained and communicated between ASes. Different from the general information such as routing information from the RIB, there are no existing mechanism which can support the perception and communication of SAV-specific information between ASes. Hence, a SAV-specific communication mechanism is needed to provide a medium and set of rules to establish communication between different ASes for the exchange of SAV-specific information.

Furthermore, in order to obtain all the source prefixes of an AS, the inter-domain SAVNET architecture may communicate with the intra-domain SAVNET architecture [intra-domain-arch] to obtain all the prefixes belonging to an AS. If the legitimate incoming directions of SAV-specific information are learned from BGP AS_PATH information, it needs to note that BGP AS_PATH may hide some interfaces that exist between ASes in the scenarios involving topological fork and merge.

Some scenarios, such as the ones where policy-based routing or static route exist in the inter-domain networks, may rely on the wider deployment of SAVNET agent or more interactive communication between ASes to make the inter-domain SAVNET work better. In these scenarios, operators may override the default BGP decision by using policy-based routing or static route. For example, in Figure 8, AS 2 may use another AS which does not in the AS path [AS 1, AS 2, AS 3] to transmit the legitimate traffic with source addresses in P1 to AS 3. For such cases, inter-domain SAVNET may require AS 2 to deploy SAVNET agent to obtain the SAV-specific information for the legitimate traffic with source addresses in P1, or AS 1 and AS 3 may need more interactive communication to obtain the SAV-specific information.

The preferred AS paths of an AS may change over time due to route changes caused by operator configurations or network failures. In addition, the SAVNET agent should be aware of the route changes and launch SAV-specific messages to adapt to the route changes in a timely manner. The SAV-specific information communication mechanism should handle route changes carefully to avoid improper blocks. The reasons for leading to improper blocks may include late detection of route changes, delayed message transmission, or packet losses. If the SAVNET agent cannot be aware of the route changes caused by failures, it may not be aware of the failures. A wider deployment of SAVNET agent can make network failure sensing more sensitive. However, the detailed design of SAV-specific information communication mechanism for dealing with route changes is outside the scope of this document.

7.2. General Information Communication Mechanism

The general information communication mechanism is used for communicating routing information between ASes, obtaining RPKI ROA objects and ASPA objects from RPKI cache servers, and obtaining the information about ASes and their prefixes from IRR databases. The general communication mechanism can be implemented by using existing protocols for collecting the relative information, such as BGP, RTR [RFC8210], and FTP [RFC959].

7.3. Management Information Communication Mechanism

The primary purpose of the management information communication mechanism is to deliver manual configurations of network operators. Examples of the management configurations include, but are not limited to:

- * SAVNET configurations using YANG, CLI, RTBH, or Flowspec: Inter-domain SAVNET implementations on vendor devices need to accept the configurations from operators, such as the SAV rules directly configured by operators, and may support various tools to do this, such as YANG, CLI, RTBH, or Flowspec.
- * SAVNET performance analysis: Inter-domain SAVNET implementations need to support various operation methods to report the statistics of SAVNET, such as alarm and exception reporting, performance monitoring and reporting for the control plane and data plane, which are discussed in detail in Section 10.
- * SAVNET deployment provisioning: Inter-domain SAVNET implementations may support the configurations which relate to its deployment process, such as maximum hardware resources used, "on-off" of SAVNET, and access authority.

Note that the management information can be delivered at any time and requires reliable delivery for the management information communication mechanism implementation. Additionally, to support performance analysis, the management information communication mechanism can carry telemetry information, such as metrics pertaining to forwarding performance, the count of spoofing packets and discarded packets, and the information regarding the prefixes associated with the spoofing traffic, as observed until the most recent time.

8. Use Cases

This section utilizes the sample use cases to showcase that the inter-domain SAVNET architecture can improve the validation accuracy in the scenarios of limited propagation of prefixes, hidden prefixes, reflection attacks, and direct attacks, compared to existing SAV mechanisms, which are also utilized for the gap analysis of existing inter-domain SAV mechanisms in [inter-domain-ps]. In the following, these use cases are discussed for SAV at customer interfaces and SAV at provider/peer interfaces, respectively.

8.1. SAV at Customer Interfaces

In order to prevent the source address spoofing, operators can enable ACL-based ingress filtering, source-based RTBH filtering, and/or uRPF-based mechanisms at customer interfaces, namely Strict uRPF, FP-uRPF, VRF uRPF, or EFP-uRPF [manrs] [nist]. However, as analyzed in [inter-domain-ps], uRPF-based mechanisms may lead to false positives in two inter-domain scenarios: limited propagation of prefixes and hidden prefixes, or may lead to false negatives in the scenarios of source address spoofing attacks within a customer cone, while ACL-based ingress filtering and source-based RTBH filtering need to update SAV rules in a timely manner and lead to high operational overhead. The following showcases that the inter-domain SAVNET architecture can avoid false positives and false negatives in these scenarios.

8.1.1. Limited Propagation of Prefixes

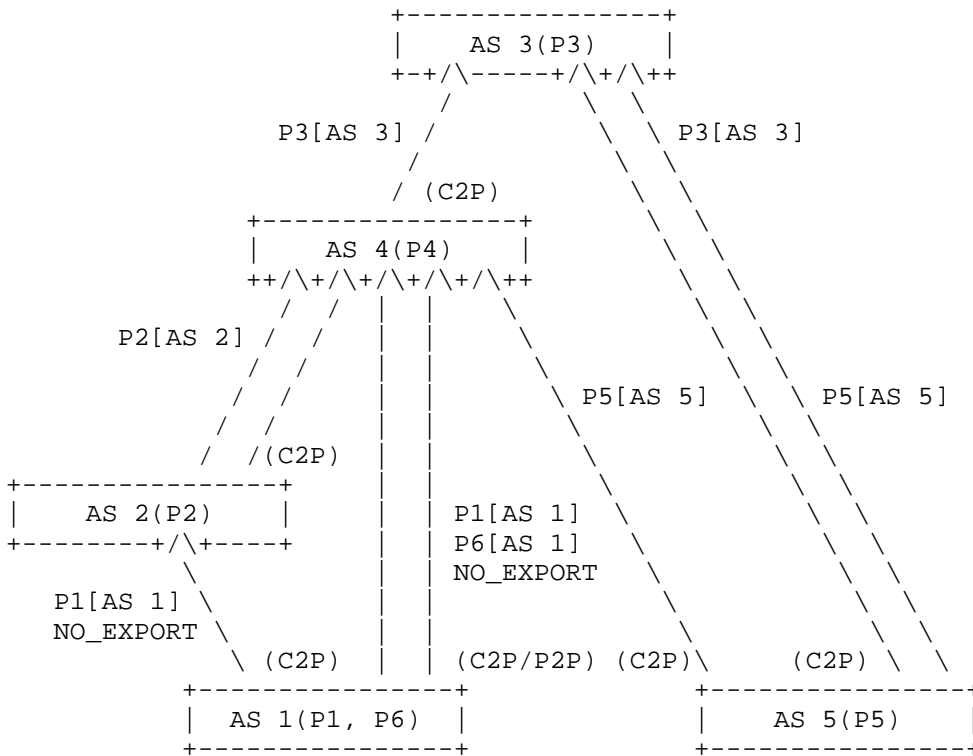


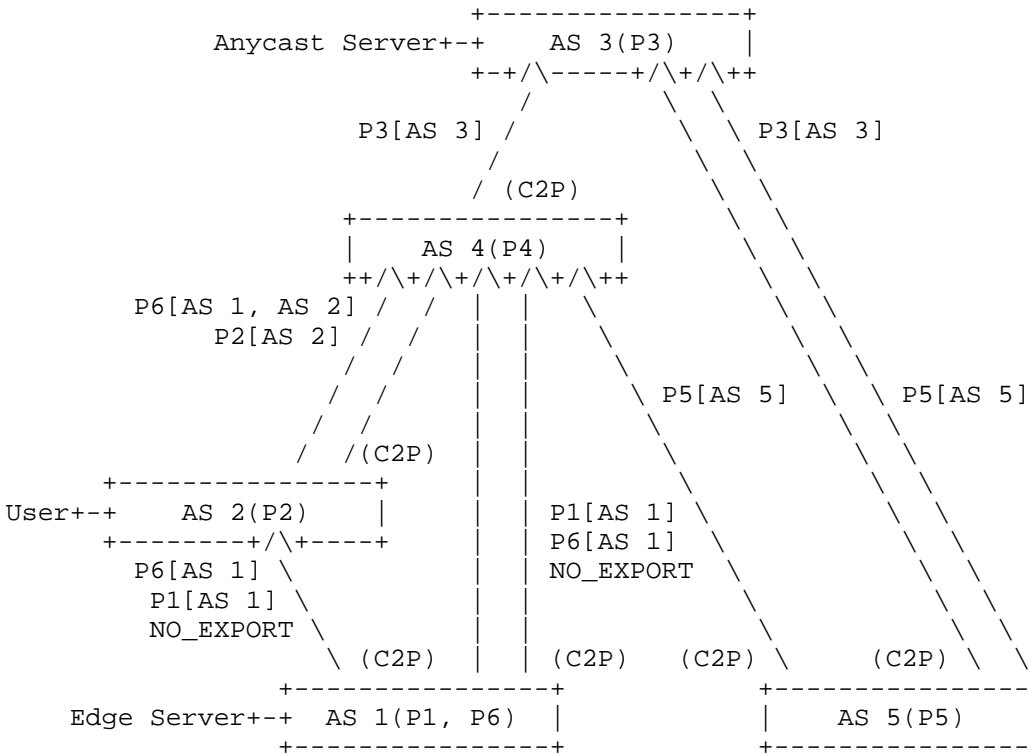
Figure 9: Limited propagation of prefixes caused by NO_EXPORT.

Figure 9 presents a scenario where the limited propagation of prefixes occurs due to the NO_EXPORT community attribute. In this scenario, AS 1 is a customer of AS 2, AS 2 is a customer of AS 4, AS 4 is a customer of AS 3, and AS 5 is a customer of both AS 3 and AS 4. The relationship between AS 1 and AS 4 can be either customer-to-provider (C2P) or peer-to-peer (P2P). AS 1 advertises prefixes P1 to AS 2 and adds the NO_EXPORT community attribute to the BGP advertisement sent to AS 2, preventing AS 2 from further propagating the route for prefix P1 to AS 4. Similarly, AS 1 adds the NO_EXPORT community attribute to the BGP advertisement sent to AS 4, resulting in AS 4 not propagating the route for prefix P6 to AS 3. Consequently, AS 4 only learns the route for prefix P1 from AS 1 in this scenario. Suppose AS 1 and AS 4 have deployed inter-domain SAV while other ASes have not, and AS 4 has deployed EFP-uRPF at its customer interfaces.

In this scenario, existing uRPF-based SAV mechanisms would block the traffic with P1 as source addresses improperly, and thus suffer from the problem of false positives [inter-domain-ps]. If the inter-

domain SAVNET architecture is deployed, AS 1 can communicate the SAV-specific information to AS 4 and AS 4 will be aware that the traffic with P1 as source addresses can arrive at the interfaces facing AS 1 and AS 2. As a result, the false positive problem can be avoided.

8.1.2. Hidden Prefixes



P3 is the anycast prefix and is only advertised by AS 3 through BGP.

Figure 10: A Direct Server Return (DSR) scenario.

Figure 10 illustrates a direct server return (DSR) scenario where the anycast IP prefix P3 is only advertised by AS 3 through BGP. In this example, AS 3 is the provider of AS 4 and AS 5, AS 4 is the provider of AS 1, AS 2, and AS 5, and AS 2 is the provider of AS 1. AS 1 and AS 4 have deployed inter-domain SAV, while other ASes have not. When users in AS 2 send requests to the anycast destination IP, the forwarding path is AS 2->AS 4->AS 3. The anycast servers in AS 3 receive the requests and tunnel them to the edge servers in AS 1. Finally, the edge servers send the content to the users with source addresses in prefix P3. The reverse forwarding path is AS 1->AS 4->AS 2.

8.1.3. Reflection Attacks

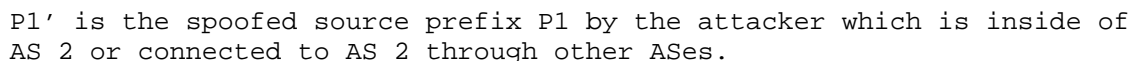
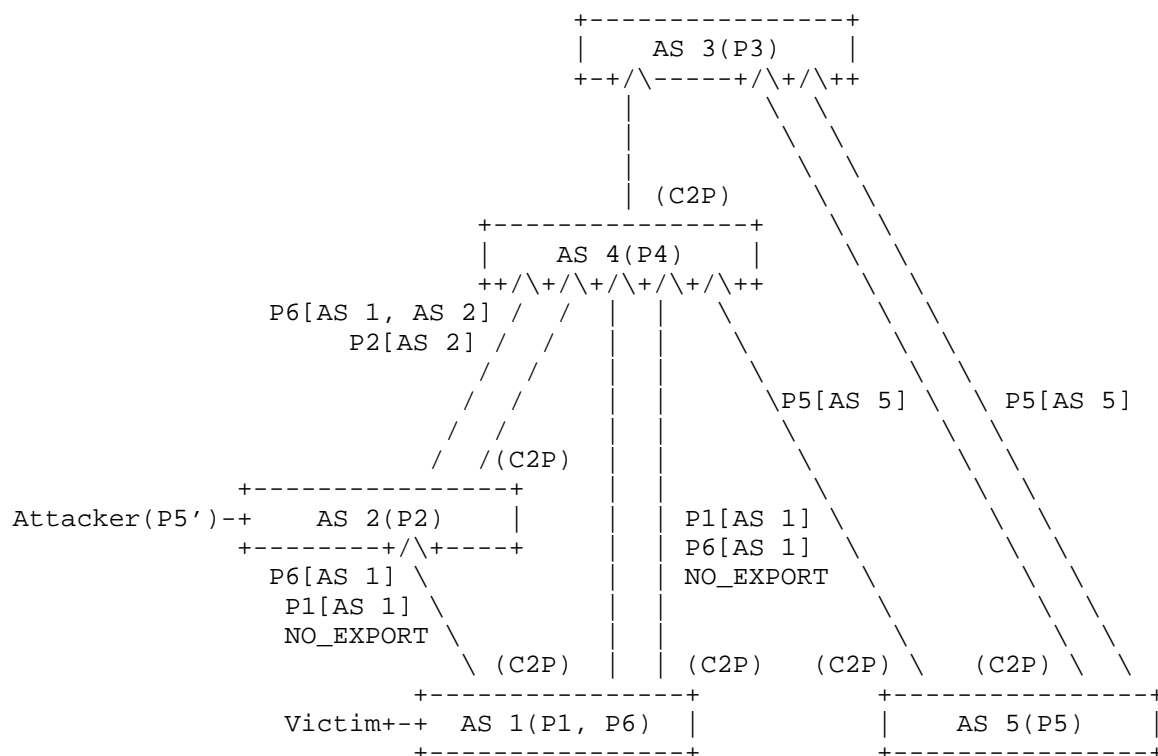


Figure 11 depicts the scenario of reflection attacks by source address spoofing within a customer cone. The reflection attack by source address spoofing takes place within AS 4's customer cone, where the attacker spoofs the victim's IP address (P1) and sends

In this scenario, EFP-uRPF with algorithm A/B will improperly permit the spoofing attacks originating from AS 2 [inter-domain-ps]. If the inter-domain SAVNET architecture is deployed, AS 1 can communicate the SAV-specific information to AS 4 and AS 4 will be aware that the traffic with P1 as source addresses can only arrive at the interface facing AS 1. Therefore, at the interface of AS 4 facing AS 2, the spoofing traffic can be blocked.

8.1.4. Direct Attacks



P1' is the spoofed source prefix P1 by the attacker which is inside of AS 2 or connected to AS 2 through other ASes.

Figure 12: A scenario of the direct attacks by source address spoofing within a customer cone.

Figure 12 portrays a scenario of direct attacks by source address spoofing within a customer cone and is used to analyze the gaps of uRPF-based mechanisms below. The direct attack by source address spoofing takes place within AS 4's customer cone, where the attacker spoofs a source address (P5) and directly targets the victim's IP address (P1), overwhelming its network resources. The arrows in Figure 12 illustrate the commercial relationships between ASes. AS 3 serves as the provider for AS 4 and AS 5, while AS 4 acts as the provider for AS 1, AS 2, and AS 5. Additionally, AS 2 is the provider for AS 1. Suppose AS 4 and AS 5 have deployed inter-domain SAV, while the other ASes have not.

In this scenario, EFP-uRPF with algorithm A/B will improperly permit the spoofing attacks [inter-domain-ps]. If the inter-domain SAVNET architecture is deployed, AS 5 can communicate the SAV-specific information to AS 4 and AS 4 will be aware that the traffic with P5 as source addresses can arrive at the interface facing AS 3 and AS 5. Therefore, at the interface of AS 4 facing AS 2, the spoofing traffic can be blocked.

8.2. SAV at Provider/Peer Interfaces

In order to prevent packets with spoofed source addresses from the provider/peer AS, ACL-based ingress filtering, Loose uRPF, and/or source-based RTBH filtering can be deployed [nist]. [inter-domain-ps] exposes the limitations of ACL-based ingress filtering, source-based RTBH filtering, and Loose uRPF for SAV at provider/peer interfaces in scenarios of source address spoofing attacks from provider/peer AS. The source address spoofing attacks from provider/peer AS include reflection attacks from provider/peer AS and direct attacks from provider/peer AS. The following showcases that the inter-domain SAVNET architecture can avoid false negatives in these scenarios.

8.2.1. Reflection Attacks

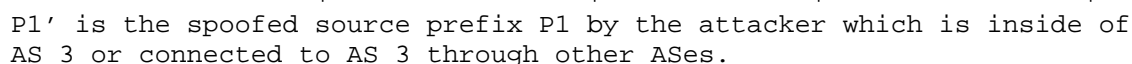
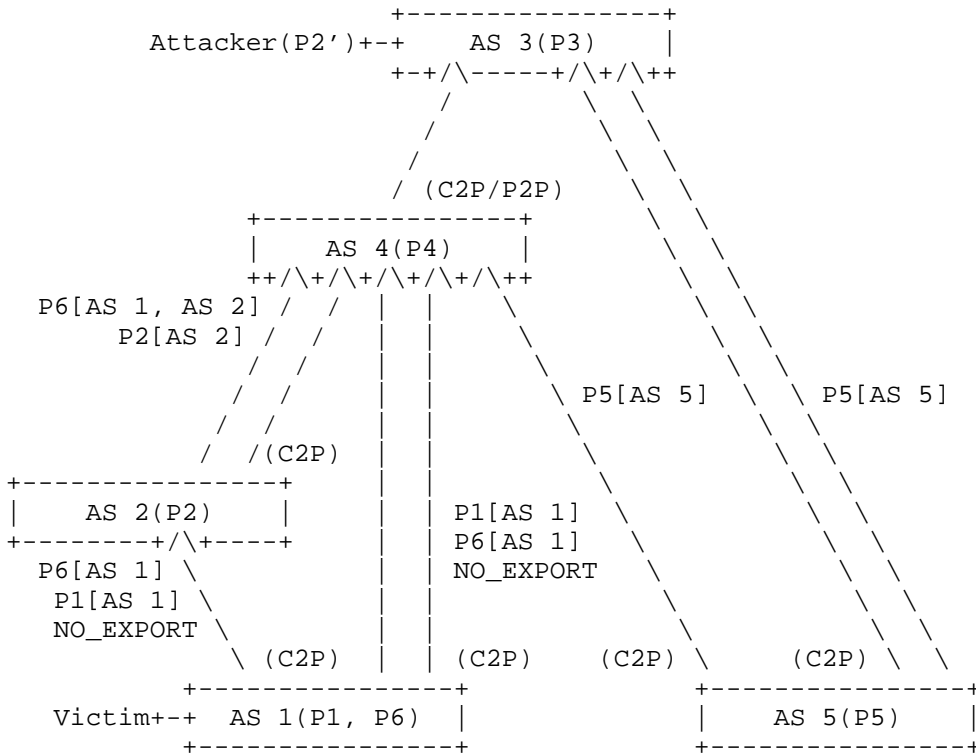


Figure 13 depicts the scenario of reflection attacks by source address spoofing from provider/peer AS. In this case, the attacker spoofs the victim's IP address (P1) and sends requests to servers' IP address (P2) that respond to such requests. The servers then send overwhelming responses back to the victim, exhausting its network resources. The arrows in Figure 13 represent the commercial relationships between ASes. AS 3 acts as the provider or lateral peer of AS 4 and the provider for AS 5, while AS 4 serves as the provider for AS 1, AS 2, and AS 5. Additionally, AS 2 is the provider for AS 1. Suppose AS 1 and AS 4 have deployed inter-domain SAV, while the other ASes have not.

Both ACL-based ingress filtering and source-based RTBH filtering will induce additional operational overhead, and Loose uRPF may improperly permit spoofed packets [inter-domain-ps]. If the inter-domain SAVNET architecture is deployed, AS 1 can communicate the SAV-specific

information to AS 4 and AS 4 will be aware that the traffic with P1 as source addresses can arrive at the interface facing AS 1 and AS 2. Therefore, at the interface of AS 4 facing AS 3, the spoofing traffic can be blocked.

8.2.2. Direct Attacks



P2' is the spoofed source prefix P2 by the attacker which is inside of AS 3 or connected to AS 3 through other ASes.

Figure 14: A scenario of direct attacks by source address spoofing from provider/peer AS.

Figure 14 showcases a scenario of direct attack by source address spoofing from provider/peer AS. In this case, the attacker spoofs another source address (P2) and directly targets the victim's IP address (P1), overwhelming its network resources. The arrows in Figure 14 represent the commercial relationships between ASes. AS 3 acts as the provider or lateral peer of AS 4 and the provider for AS 5, while AS 4 serves as the provider for AS 1, AS 2, and AS 5. Additionally, AS 2 is the provider for AS 1. Suppose AS 1 and AS 4 have deployed inter-domain SAV, while the other ASes have not.

Also, in this scenario, both ACL-based ingress filtering and source-based RTBH filtering will induce additional operational overhead, and Loose uRPF may improperly permit spoofed packets [inter-domain-ps]. If the inter-domain SAVNET architecture is deployed, AS 2 can communicate the SAV-specific information to AS 4 and AS 4 will be aware that the traffic with P2 as source addresses can only arrive at the interface facing AS 2. Therefore, at the interface of AS 4 facing AS 3, the spoofing traffic can be blocked.

9. Meeting the Design Requirements of Inter-domain SAVNET

The inter-domain SAVNET architecture proposes the guidelines for the design of new inter-domain SAV mechanisms to meet the requirements defined in [inter-domain-ps]. The followings illustrate the design guidelines to meet the requirements one by one.

9.1. Improving Validation Accuracy over Existing Mechanisms

As analyzed in Section 8, existing uRPF-based SAV mechanisms may have improper block or improper permit problems in the scenarios of limited propagation of prefixes, hidden prefixes, reflection attacks, and direct attacks for SAV at customer interfaces, and the scenarios of reflection attacks and direct attacks for SAV at provider/peer interfaces.

Inter-domain SAVNET proposes SAV-specific information, which consists of the source prefixes of ASes and their corresponding legitimate incoming direction to enter other ASes. ASes which deploy SAVNET agent can communicate SAV-specific information with each other and generate accurate SAV rules for the prefixes from the SAV-specific information. The use cases shown in Section 8 has demonstrated inter-domain SAVNET can improve validation accuracy compared to existing SAV mechanisms. Along with more ASes deploy SAVNET agent and communicate SAV-specific information with each other, accurate SAV rules can be generated for these ASes and their prefixes can obtain better protection.

9.2. Working in Incremental/Partial Deployment

A new inter-domain SAVNET mechanism should consider incremental/partial deployment as it is not feasible to deploy SAVNET agent simultaneously in all ASes, due to various constraints, such as device capabilities, versions, or vendors.

Inter-domain SAVNET can support incremental/partial deployment, as it is not mandatory for all ASes to deploy SAVNET agents for communicating SAV-specific information. ASes which deploy SAVNET agents can establish a logical neighboring relationship with other

ASes. The connections for communicating SAV-specific information can be achieved by manual configurations set by operators or an automatic neighbor discovery mechanism. An automatic neighbor discovery mechanism can utilize existing protocols or tools to collect the SAVNET neighboring information. This flexibility enables the inter-domain SAVNET to accommodate varying degrees of deployment, promoting interoperability and collaboration among participating ASes. During the partial/incremental deployment of SAVNET agent, the SAV-specific information for the ASes which do not deploy SAVNET agent cannot be obtained. To protect the prefixes of these ASes, inter-domain SAVNET can use the general information in the SIB to generate SAV rules. When using the general information, inter-domain SAVNET needs to guarantee the SAV accuracy for the corresponding application scenarios. The use cases in Section 8 demonstrates that inter-domain SAVNET supports incremental/partial deployment.

As more ASes adopt the inter-domain SAVNET, the "deployed area" expands, thereby increasing the collective defense capability against source address spoofing. Furthermore, if multiple "deployed areas" can be logically interconnected across "non-deployed areas", these interconnected "deployed areas" can form a logical alliance, providing enhanced protection against address spoofing. Especially, along with more ASes deploy SAVNET agent and support the communication of SAV-specific information, the generated SAV rules will become more accurate, as well as enhancing the protection capability against source address spoofing.

In addition, releasing the SAV functions of the inter-domain SAVNET incrementally is RECOMMENDED as one potential way to reduce the deployment risks and can be considered in its deployment by network operators:

- * First, the inter-domain SAVNET can only do the measurement in the data plane and do not take any other actions. Based on the measurement data, the operators can evaluate the effect of the inter-domain SAVNET on the legitimate traffic, including validation accuracy and forwarding performance, as well as the operational overhead.
- * Second, the inter-domain SAVNET can open the function to limit the rate of the traffic that is justified as spoofing traffic. The operators can further evaluate the effect of the inter-domain SAVNET on the legitimate traffic and spoofing traffic, such as limiting the rate of all the spoofing traffic without hurting the legitimate traffic.

- * Third, when the validation accuracy, forwarding performance, and operational overhead have been verified on a large scale by the live network, the inter-domain SAVNET can open the function to directly block the spoofing traffic that is justified by the SAV table in the data plane.

9.3. Reducing Operational Overhead

The inter-domain routes or the prefixes of ASes usually change dynamically, which requires the SAV rules to be updated automatically. ACL-based ingress filtering and source-based RTBH filtering requires manual configuration to update SAV rules to adapt to the routing or prefix changes, which leads to high operational overhead.

Inter-domain SAVNET proposes the SAV-specific information communication mechanism and utilizes it to communicate SAV-specific information automatically between ASes which deploy SAVNET agent. Upon receiving the SAV-specific information, SAVNET agent will use it to generate SAV rules. The use cases displayed in Section 8.2 show that inter-domain SAVNET reduces operational overhead compared to ACL-based ingress filtering and source-based RTBH filtering.

9.4. Guaranteeing Convergence

Convergence issues SHOULD be carefully considered due to the dynamic nature of the Internet. Internet routes undergo continuous changes, and SAV rules MUST proactively adapt to these changes, such as prefix and route changes, in order to avoid improper block and reduce improper permit. To effectively track these changes, the SAVNET agent should proactively communicate the changes of SAV-specific information between ASes and generate SAV rules in a timely manner.

The SAVNET agent should launch SAV-specific messages to adapt to the route or prefix changes in a timely manner. During the routing convergence process, the traffic paths of the source prefixes can undergo rapid changes within a short period. The changes of the SAV-specific information may not be communicated in time between ASes to update SAV rules, improper block or improper permit may happen. Such inaccurate validation is caused by the delays in communicating SAV-specific information between ASes, which occur due to the factors like packet losses, unpredictable network latencies, or message processing latencies. The detailed design of the SAV-specific information communication mechanism should consider these issues to reduce the inaccurate validation.

Besides, for the inter-domain SAVNET, the potential ways to deal with the inaccurate validation issues during the convergence of the SAV-specific information communication mechanism is to consider using the information from RPKI ROA objects and ASPA objects to generate SAV rules until the convergence process of the SAV-specific communication mechanism is finished, since these information is more stable and can help avoid improper block, and thus avoiding the impact to the legitimate traffic.

9.5. Providing Necessary Security Guarantee

For inter-domain SAVNET, the SAVNET agent plays a crucial role in generating and disseminating SAV-specific messages across different ASes. To safeguard against the potential risks posed by a malicious AS generating incorrect or forged SAV-specific messages, it is important for the SAVNET agents to employ security authentication measures for each received SAV-specific message. The major security threats faced by inter-domain SAVNET can be categorized into two aspects: session security and content security. Session security pertains to verifying the identities of both parties involved in a session and ensuring the integrity of the session content. Content security, on the other hand, focuses on verifying the authenticity and reliability of the session content, thereby enabling the identification of forged SAV-specific messages.

The threats to session security include:

- * Session identity impersonation: This occurs when a malicious router deceitfully poses as a legitimate peer router to establish a session with the targeted router. By impersonating another router, the malicious entity can gain unauthorized access and potentially manipulate or disrupt the communication between the legitimate routers.
- * Session integrity destruction: In this scenario, a malicious intermediate router situated between two peering routers intentionally tampers with or destroys the content of the relayed SAV-specific message. By interfering with the integrity of the session content, the attacker can disrupt the reliable transmission of information, potentially leading to miscommunication or inaccurate SAV-related data being propagated.

The threats to content security include:

- * Message alteration: A malicious router has the ability to manipulate or forge any portion of a SAV-specific message. For example, the attacker may employ techniques such as using a spoofed Autonomous System Number (ASN) or modifying the AS path

information within the message. By tampering with the content, the attacker can potentially introduce inaccuracies or deceive the receiving ASes, compromising the integrity and reliability of the SAV-related information.

- * Message injection: A malicious router injects a seemingly "legitimate" SAV-specific message into the communication stream and directs it to the corresponding next-hop AS. This type of attack can be likened to a replay attack, where the attacker attempts to retransmit previously captured or fabricated messages to manipulate the behavior or decisions of the receiving ASes. The injected message may contain malicious instructions or false information, leading to incorrect SAV rule generation or improper validation.
- * Path deviation: A malicious router intentionally diverts a SAV-specific message to an incorrect next-hop AS, contrary to the expected path defined by the AS path. By deviating from the intended routing path, the attacker can disrupt the proper dissemination of SAV-related information and introduce inconsistencies or conflicts in the validation process. This can undermine the effectiveness and accuracy of source address validation within the inter-domain SAVNET architecture.

Overall, inter-domain SAVNET shares similar security threats with BGP and can leverage existing BGP security mechanisms to enhance both session and content security. Session security can be enhanced by employing session authentication mechanisms used in BGP. Similarly, content security can benefit from the deployment of existing BGP security mechanisms like RPKI, BGPsec, and ASPA. While these mechanisms can address content security threats, their widespread deployment is crucial. Until then, it is necessary to develop an independent security mechanism specifically designed for inter-domain SAVNET. One potential approach is for each source AS to calculate a digital signature for each AS path and include these digital signatures within the SAV-specific messages. Upon receiving a SAV-specific message, the SAVNET agent can verify the digital signature to ascertain the message's authenticity. Furthermore, it is worth noting that the SAV-specific information communication mechanism may need to operate over a network link that is currently under a source address spoofing attack. As a result, it may experience severe packet loss and high latency due to the ongoing attack, and the implementation of the SAV-specific communication mechanism should ensure uninterrupted communication. Detailed security designs and considerations will be addressed in a separate draft, ensuring the robust security of inter-domain SAVNET.

10. Manageability Considerations

It is crucial to consider the operations and management aspects of SAV information sources, the SAV-specific communication mechanism, SIB, SIM, and SAV table in the inter-domain SAVNET architecture. The following guidelines should be followed for their effective management:

First, management interoperability should be supported across devices from different vendors or different releases of the same product, based on a unified data model such as YANG [RFC6020]. This is essential because the Internet comprises devices from various vendors and different product releases that coexist simultaneously.

Second, scalable operation and management methods such as NETCONF [RFC6241] and syslog protocol [RFC5424] should be supported. This is important as an AS may have hundreds or thousands of border routers that require efficient operation and management.

Third, management operations, including default initial configuration, alarm and exception reporting, logging, performance monitoring and reporting for the control plane and data plane, as well as debugging, should be designed and implemented in the protocols or protocol extensions. These operations can be performed either locally or remotely, based on the operational requirements.

By adhering to these rules, the management of SAV information sources and related components can be effectively carried out, ensuring interoperability, scalability, and efficient operations and management of the inter-domain SAVNET architecture.

11. Privacy Considerations

TBD

12. IANA Considerations

This document has no IANA requirements.

13. Scope and Assumptions

In this architecture, the choice of protocols used for communication between the SIM and different SAV information sources is not limited. The inter-domain SAVNET architecture presents considerations on how to consolidate SAV-related information from various sources to generate SAV rules and perform SAV using the SAV table in the dataplane. The detailed design and implementation for SAV rule generation and SAV execution depend on the specific inter-domain SAV

mechanisms employed.

This document does not cover administrative or business agreements that may be established between the involved inter-domain SAVNET parties. These considerations are beyond the scope of this document. However, it is assumed that authentication and authorization mechanisms can be implemented to ensure that only authorized ASes can communicate SAV-related information.

This document makes the following assumptions:

- * All ASes where the inter-domain SAVNET is deployed are assumed to provide the necessary connectivity between SAVNET agent and any intermediate network elements. However, the architecture does not impose any specific limitations on the form or nature of this connectivity.
- * Congestion and resource exhaustion can occur at various points in the inter-domain networks. Hence, in general, network conditions should be assumed to be hostile. The inter-domain SAVNET architecture must be capable of functioning reliably under all circumstances, including scenarios where the paths for delivering SAV-related information are severely impaired. It is crucial to design the inter-domain SAVNET system with a high level of resilience, particularly under extremely hostile network conditions. The architecture should ensure uninterrupted communication between inter-domain SAVNET agents, even when data-plane traffic saturates the link.
- * The inter-domain SAVNET architecture does not impose rigid requirements for the SAV information sources that can be used to generate SAV rules. Similarly, it does not dictate strict rules on how to utilize the SAV-related information from diverse sources or perform SAV in the dataplane. Network operators have the flexibility to choose their approaches to generate SAV rules and perform SAV based on their specific requirements and preferences. Operators can either follow the recommendations outlined in the inter-domain SAVNET architecture or manually specify the rules for governing the use of SAV-related information, the generation of SAV rules, and the execution of SAV in the dataplane.
- * The inter-domain SAVNET architecture does not impose restrictions on the selection of the local AS with which AS to communicate SAV-specific Information. The ASes have the flexibility to establish connections for SAV-specific communication based on the manual configurations set by operators or other automatic mechanisms.

- * The inter-domain SAVNET architecture provides the flexibility to accommodate Quality-of-Service (QoS) policy agreements between SAVNET-enabled ASes or local QoS prioritization measures, but it does not make assumptions about their presence. These agreements or prioritization efforts are aimed at ensuring the reliable delivery of SAV-specific Information between SAVNET agents. It is important to note that QoS is considered as an operational consideration rather than a functional component of the inter-domain SAVNET architecture.
- * The SAVNET communication mechanisms are loosely coupled and are used for communicating or gathering SAV-related information, and how the inter-domain SAVNET synchronizes the management and operation configurations is out of scope of this document.

14. Contributors

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

Igor Lubashev
Akamai Technologies
145 Broadway
Cambridge, MA, 02142
United States of America
Email: ilubashe@akamai.com

Many thanks to Mingqing Huang and Igor Lubashev for the significantly helpful discussions and revision suggestions.

15. References

15.1. Normative References

- [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/rfc/rfc8174>>.
- [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/rfc/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/rfc/rfc8704>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/rfc/rfc6020>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/rfc/rfc6241>>.
- [RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, DOI 10.17487/RFC5424, March 2009, <<https://www.rfc-editor.org/rfc/rfc5424>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/rfc/rfc5280>>.
- [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/rfc/rfc2119>>.

15.2. Informative References

- [inter-domain-ps] "Source Address Validation in Inter-domain Networks Gap Analysis, Problem Statement, and Requirements", 2025, <<https://datatracker.ietf.org/doc/draft-ietf-savnet-inter-domain-problem-statement/>>.
- [intra-domain-arch] "Intra-domain Source Address Validation (SAVNET) Architecture", 2025, <<https://datatracker.ietf.org/doc/draft-ietf-savnet-intra-domain-architecture/>>.
- [RFC5635] Kumari, W. and D. McPherson, "Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF)", RFC 5635, DOI 10.17487/RFC5635, August 2009, <<https://www.rfc-editor.org/rfc/rfc5635>>.

- [RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <<https://www.rfc-editor.org/rfc/rfc8955>>.
- [RFC8210] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1", RFC 8210, DOI 10.17487/RFC8210, September 2017, <<https://www.rfc-editor.org/rfc/rfc8210>>.
- [RFC959] Postel, J. and J. Reynolds, "File Transfer Protocol", STD 9, RFC 959, DOI 10.17487/RFC0959, October 1985, <<https://www.rfc-editor.org/rfc/rfc959>>.
- [manrs] MANRS, "MANRS Implementation Guide", 2023, <<https://www.manrs.org/netops/guide/antispoofing/>>.
- [nist] NIST, "Resilient Interdomain Traffic Exchange: BGP Security and DDos Mitigation", 2019, <<https://www.nist.gov/publications/resilient-interdomain-traffic-exchange-bgp-security-and-ddos-mitigation>>.
- [rpki-time-of-flight] ISOC, "RPKI Time-of-Flight Tracking Delays in the Management, Control, and Data Planes", n.d., <https://dl.acm.org/doi/10.1007/978-3-031-28486-1_18>.
- [sav-table] "General Source Address Validation Capabilities", 2025, <<https://datatracker.ietf.org/doc/draft-ietf-savnet-general-sav-capabilities/>>.

Acknowledgements

Many thanks to Alvaro Retana, Kotikalapudi Sriram, Rüdiger Volk, Xueyan Song, Ben Maddison, Jared Mauch, Joel Halpern, Aijun Wang, Jeffrey Haas, Xiangqing Chang, Changwang Lin, Mingxing Liu, Zhen Tan, Yuanyuan Zhang, Yangyang Wang, Antoin Verschuren, Olaf Struck, Siyuan Teng, Gert Doering etc. for their valuable comments on this document.

Authors' Addresses

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

Li Chen
Zhongguancun Laboratory
Beijing
China
Email: lichen@zgclab.edu.cn

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

Libin Liu
Zhongguancun Laboratory
Beijing
China
Email: liulb@zgclab.edu.cn

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