

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2026

G. Ren
S. Liu
X. Yin
M.L. Jia
Tsinghua University
2 March 2026

Inter-domain Source Address Validation based on AS relationships
draft-rly-savnet-inter-domain-as-relationships-06

Abstract

This draft introduces a distributed inter-domain source address validation scheme based on AS relationships named AS Relationship Based Inter-domain Filtering (ARBIF). It primarily describes this method from seven aspects: a brief introduction to the scheme, an overview of the AS relationship classification and acquisition methods, the architecture of the ARBIF system, implementation based on BGP extension, typical use cases, experiments of ARBIF, and considerations for deployability.

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 3 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. Requirements Language	4
2. Terminology	4
3. Introduction to AS Relationships	5
3.1. Major AS relationships	6
3.2. Complex AS Relationships	8
3.3. AS relationship acquisition methods	9
3.3.1. Inference Algorithms	9
3.3.2. Querying approach	11
4. Architecture of AS Relationship Based Inter-domain Filtering (ARBIF)	11
4.1. Overall Architecture	11
4.2. Validation Router (VR)	12
4.2.1. VR Role in the ARBIF system	12
4.2.2. VR Implementation	13
4.3. AS-IP Prefix Mapping Server (AIMS)	13
4.3.1. AIMS Role in the ARBIF system	13
4.3.2. AIMS Implementation based on RPKI	14
4.3.3. Lightweight AIMS Implementation without RPKI	15
5. BGP Extension for Inter-domain SAV	16
5.1. Feasibility of BGP Extension	16
5.2. Implementation of BGP Extension	16
5.3. An example of BGP Extension	17
6. Scenarios	18
6.1. Multi-homing Scenarios	18
6.1.1. Multipoint Interconnection Scenario	18
6.1.2. Multi-homing Scenario	19
6.2. Dynamic Scenario	20
6.2.1. AS Relationships Change	21
6.2.2. AS Prefixes Change	21
6.2.3. AS Network Topologies Change	22
6.2.4. BGP Attributes Change	22
6.3. IXP Scenario	23
7. Experiment of ARBIF Implementation	24
7.1. Environment	24
7.2. Implementation Method	25
7.3. Simulation Experiment Result	25
8. Considerations on Deployability	26
8.1. Utilize existing information as much as possible	26
8.2. Prefer to use and exchange more abstract information	26
8.3. Balance accuracy, time and cost	26

9. Next Step	27
10. Security Considerations	27
11. IANA Considerations	27
12. References	27
12.1. Normative References	27
12.2. Informative References	28
Acknowledgements	30
Authors' Addresses	30

1. Introduction

Various attacks continue to pose significant security threats to the Internet, and IP spoofing is critical. Attackers frequently employ IP spoofing to launch DDoS attacks and disguise their actual identities. Source address validation (SAV) can greatly relieve IP spoofing and mitigate DDoS attacks.

The Source Address Validation Architecture (short for SAVA) proposed by [RFC5210] divides its SAV architecture into three levels: the access network, intra-domain, and inter-domain. In SAV at the access network level, many researchers have made considerable progress and established several standards through discussion and collaboration.

Researchers also proposed algorithms for inter-domain SAV. [RFC2827], [RFC3704], and [RFC8704] proposed uRPF algorithms that reverse routers' forwarding tables as their SAV rules. They further proposed several variants based on this core idea to fit different scenarios. uRPF algorithms exhibit high convergence speed and low cost. The SAVNET working group is devoted to improving the inter-domain SAV mechanism [inter-domain-sav-ps] and designing an SAV architecture using various information [inter-domain-sav-archt]. Their scheme exhibits high accuracy. The BAR-SAV algorithm [sidrops-bar-sav] in the SIDRops working group generates a permissible prefix list as SAV rules using BGP UPDATE messages, ASPA, and ROA objects in the RPKI. It has medium accuracy.

Though all existing methods have advantages, they have yet to become an effective and deployable standard. Aiming to fix this gap, we propose a scheme with moderate accuracy, convergence speed, and cost. To implement it, we use AS relationships to abstract the inter-domain routing information.

At the AS level, each AS owns some IP address prefixes and advertises them to neighbor ASes. Through its advertisement, neighbor ASes know they can route traffic to these prefixes through it. What's more, neighbor ASes also determine whether to propagate the received routes to their neighbor ASes according to AS relationships. Thus, we can estimate each prefix's propagation, and infer approximate inter-domain routes using AS relationships and IP address prefixes.

This scheme's inaccuracy comes from ignoring fine-grained routing information, such as BGP path attributes. Ignoring them may cause routes to propagate beyond the intended scope, leading to more improper permits. Even one dropped legitimate packet may lead to serious Internet interruption, but a few passed spoofed packets cannot cause large-scale attacks. As our scheme does not cause additional improper blocks, it does not violate requirements for inter-domain SAV.

This draft introduces a distributed inter-domain SAV scheme based on AS relationships, and we name it AS Relationship Based Inter-domain Filtering (ARBIF). Receiving comments from other researchers about deployment upgrade costs, the ARBIF scheme adopts a distributed SAV architecture without a centralized server in each AS or a newly designed protocol. Instead, we extend the current BGP protocol and directly implement the ARBIF on existing AS border routers. These are ARBIF's main modifications compared to the original scheme in [RFC5210]. ARBIF also covers more AS relationships and discusses more scenarios. We will explain its details in the following sections.

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 AS or source IP prefix.

ASN SAV Rule:

The rule that indicates the validity of specific source ASes and is usually in the form of an AS number (ASN) set.

IP Prefix SAV Rule:

The rule that indicates the validity of specific source IP prefixes and is usually in the form of an IP prefix set.

Validation Router (VR):

The border router of a specific AS in the ARBIF system that takes responsibility for exchanging and generating ASN SAV rules, generating IP Prefix SAV rules using the mapping from AS numbers to IP prefixes, and validating packets.

AS-IP Prefix Mapping Server (AIMS):

The only centralized server in the ARBIF system that takes responsibility for maintaining the mapping from ASN to IP prefixes and providing this mapping for VRs to generate IP Prefix SAV rules according to ASN SAV rules.

Neighbor SAV Rule Table:

The table in a specific VR that records SAV Rules at all its interfaces facing neighbor ASes, including ASN SAV rules and IP Prefix SAV rules.

Improper Block:

The situation in which packets with legitimate source addresses are blocked, causing SAV false positives.

Improper Permit:

The situation in which packets with spoofed source addresses are allowed, causing SAV false negatives.

3. Introduction to AS Relationships

AS relationships are essentially business relationships between autonomous systems. Some major relationships occupy the maximal proportion of all AS relationships, while other complex relationships exist in particular situations.

To formally describe AS relationships, we define some symbols in Table 1.

Symbol	Symbol Meaning
Cus(AS_u)	Customer AS of AS_u
Pro(AS_u)	Provider AS of AS_u
Peer(AS_u)	Peer AS of AS_u
Sib(AS_u)	Sibling AS of AS_u
Hybrid(AS_u)	AS connected to AS_u in hybrid relationship
PartCus(AS_u)	Customer AS of AS_u in Partial Transit relationship
PartPro(AS_u)	Provider AS of AS_u in Partial Transit relationship
AS_uA	Position A of AS_u
RI(AS_u)	Routing Information of AS_u
EXRI(AS_1, AS_2)	Routing Information exported from AS_1 to AS_2

Table 1: Symbol definitions of formal descriptions

3.1. Major AS relationships

The major AS relationships include three different types: Provider-to-customer, Peer-to-peer, and Sibling-to-sibling relationships. The definitions and descriptions of them are as follows.

I Provider-to-customer Relationship (Transit Relationship, P2C Relationship)

The provider and customer ASes usually do not belong to the same organization. A customer AS pays its provider AS for connectivity to the rest of the Internet. Therefore, a provider AS does transit traffic for its customer ASes [infer-relatsh]. The provider AS exports all its routes to its customer because its customer pays for all traffic, while the customer AS only exports its routes, its customer routes, and its sibling routes to its provider. The formal description of the P2C relationship is as follows.

$$\text{EXRI}(\text{AS}, \text{Pro}(\text{AS})) = \text{RI}(\text{AS}) \cup \text{RI}(\text{Cus}(\text{AS})) \cup \text{RI}(\text{Sib}(\text{AS}))$$

$$\text{EXRI}(\text{AS}, \text{Cus}(\text{AS})) = \text{RI}(\text{AS}) \cup \text{RI}(\text{Cus}(\text{AS})) \cup \text{RI}(\text{Sib}(\text{AS})) \cup \text{RI}(\text{Peer}(\text{AS})) \cup \text{RI}(\text{Pro}(\text{AS}))$$

II Peer-to-peer Relationship (P2P Relationship)

A pair of peer ASes usually do not belong to the same organization but agree to exchange traffic between their customers free of charge [infer-relatsh]. Each peer AS only exports its routes, its customer routes, and its sibling routes to the other AS. The formal description of the P2P relationship is as follows.

$$\text{EXRI}(\text{AS}, \text{Peer}(\text{AS})) = \text{RI}(\text{AS}) \cup \text{RI}(\text{Cus}(\text{AS})) \cup \text{RI}(\text{Sib}(\text{AS}))$$

III Sibling-to-sibling Relationship (S2S Relationship)

Two sibling ASes are operated by the same institution. The most common anomalies stem from recent acquisitions and mergers, suggesting that some AS pairs may have a sibling relationship. Each AS exports all its routes to its sibling ASes [charact-inet]. The formal description of the S2S relationship is as follows.

$$\text{EXRI}(\text{AS}, \text{Sib}(\text{AS})) = \text{RI}(\text{AS}) \cup \text{RI}(\text{Cus}(\text{AS})) \cup \text{RI}(\text{Sib}(\text{AS})) \cup \text{RI}(\text{Peer}(\text{AS})) \cup \text{RI}(\text{Pro}(\text{AS}))$$

Based on the above descriptions of the three major AS relationships, we summarize their export rules in Table 2.

	Peer	Provider	Customer	Sibling	Self
to Peer			+	+	+
to Provider			+	+	+
to Customer	+	+	+	+	+
to Sibling	+	+	+	+	+

Table 2: Export Rule Table of Major AS Relationships

3.2. Complex AS Relationships

The major AS relationships introduced in Section 3.1 cannot cover all practical scenarios, and researchers have discovered other complex AS relationships. This draft illustrates only two relatively common ones, hybrid and partial transit relationships, as representatives. However, it is significant to note that more complex AS relationships may appear with the further development of Internet applications.

I Hybrid Relationship

Two ASes with the hybrid relationship have different relationships at different interconnection points (e.g., P2C in one location and P2P elsewhere) [inferring-complex]. According to the definition, the AS relationship between a pair of interconnection points decides their export rules.

To explain it clearly, we take the hybrid of P2C and P2P relationships as an example. We assume that AS 1 and AS 2 are in a hybrid relationship, and AS 1 is AS 2's provider at point A while they are peers at point B, as shown in Figure 1.




Figure 1: An example of Hybrid Relationships

The formal descriptions of the hybrid relationship at points A and B are as follows.

$$\text{EXRI}(\text{AS_1A}, \text{Hybrid}(\text{AS_1A})) = \text{RI}(\text{AS_1A}) \cup \text{RI}(\text{Cus}(\text{AS_1A})) \cup \text{RI}(\text{Sib}(\text{AS_1A})) \cup \text{RI}(\text{Peer}(\text{AS_1A})) \cup \text{RI}(\text{Pro}(\text{AS_1A}))$$

$$\text{EXRI}(\text{AS_1B}, \text{Hybrid}(\text{AS_1B})) = \text{RI}(\text{AS_1B}) \cup \text{RI}(\text{Cus}(\text{AS_1B})) \cup \text{RI}(\text{Sib}(\text{AS_1B}))$$

This example uses formal descriptions to display the route export rules between ASes in a hybrid relationship. The key idea is to deal with routes at different interconnection points separately.

II Partial Transit Relationship

The Partial Transit relationship restricts the scope of a typical P2C relationship to the provider AS's peer ASes and customer ASes (but not provider ASes) [inferring-complex]. According to the definition, the formal descriptions of this relationship for the AS providing partial connection services and the AS using partial connection services are as follows. And Table 3 shows export rules of partial transit relationship.

$$\text{EXRI}(\text{AS}, \text{PartCus}(\text{AS})) = \text{RI}(\text{AS}) \cup \text{RI}(\text{Cus}(\text{AS})) \cup \text{RI}(\text{Sib}(\text{AS})) \cup \text{RI}(\text{Peer}(\text{AS}))$$

$$\text{EXRI}(\text{AS}, \text{PartPro}(\text{AS})) = \text{RI}(\text{AS}) \cup \text{RI}(\text{Cus}(\text{AS})) \cup \text{RI}(\text{Sib}(\text{AS}))$$

	Peer	Provider	Customer	Sibling	Self
to Partial-Customer	+		+	+	+
to Partial-Provider			+	+	+

Table 3: Export Rule Table of Partial Transit Relationship

3.3. AS relationship acquisition methods

Several methods can obtain AS relationships with existing data, such as BGP route information, IXP information, IRR database, and ASPA objects in RPKI et al. Researchers divide these methods into two categories. One is to infer relationships between ASes using specific network data, and the other is to query data directly to obtain AS relationships.

3.3.1. Inference Algorithms

Previous researchers have proposed various AS relationship inference algorithms using different strategies.

The earliest AS relationship inferring algorithm was proposed by Gao, which speculates on AS relationships based on the Valley Free principle and observation of network phenomena [infer-relatsh]. Gao algorithm believes that the scale of provider AS is usually more immense than that of customer AS. It also supposes that the scale of one AS is generally proportional to its degree in the AS topology graph. Therefore, the Gao algorithm sorts all ASes according to

their degrees and assigns each AS connection a relationship based on the sorting results. Overall, the Gao algorithm is easy to implement and has low time complexity, but its accuracy is also low. Threshold parameters used in the algorithm will affect the inference results of AS relationships. Therefore, manual parameter selection requires much experience. Many subsequent AS relationship inferring algorithms are also based on Gao's Valley Free principle.

The AS Rank algorithm [as-rank] proposed by Luckie et al. does not rely on Gao's Valley Free principle but proposes three hypotheses as the algorithm foundation: firstly, multiple large provider ASes form peer-to-peer networks to provide global connectivity, building a set of ASes at the top of the hierarchy; Secondly, provider ASes will export its client ASes' routes to its provider ASes, and ASes outside the peer-to-peer network composed of large provider ASes need to connect with provider ASes to obtain global connectivity; Thirdly, the topological connections of AS can be represented using directed acyclic graphs. Based on the three assumptions, the AS Rank algorithm can infer P2C and P2P relationships but cannot handle other complex AS relationships. The AS Rank algorithm exhibits high accuracy and recall and can correctly infer 99.6% of P2C relationships and 98.7% of P2P relationships in validation experiments. The AS relationships inferred with the AS Rank algorithm are still continuously updated on CAIDA.

As the AS Rank algorithm has shown excessively high inferring accuracy on public datasets, the probabilistic algorithm Problink proposed by Jin et al. aims to improve the inferring accuracy in some complex situations [problink]. The Problink algorithm is based on a naive Bayesian framework and reveals crucial AS connection features derived from stochastically informative signals. Problink exhibits a lower error rate than the AS Rank algorithm on the whole dataset, especially in complex AS relationship inferring situations.

With the development and progress of AI technology, some researchers also attempted to apply advanced AI technologies to AS relationship inferring. Varghese et al. use machine learning algorithms to train one AdaBoost model for inferring AS relationships [ml-pred]. The BGP2Vec algorithm embeds ASes in a vector space for relationship classification, referring to the NLP word embedding method Word2Vec [bgp2vec]. However, these methods have relatively low accuracy and interpretability, so they do not receive much attention.

3.3.2. Querying approach

Apart from the inference algorithms in Section 3.3.1, we can also directly obtain AS relationships by querying ASPA objects in RPKI. An ASPA object is a cryptographically verifiable attestation by a Customer AS (CAS) containing a list of its authorized provider ASes [sidrops-aspa]. Therefore, we can directly get an AS's provider ASes and customer ASes from ASPA objects. Some researchers proposed the [sidrops-asra] (Autonomous System Relationship Authorization) object based on ASPA. ASRA objects can record more information about more complex AS relationships and may help us directly obtain accurate AS relationships in the future.

In this draft, as ASes in the ARBIF system make an appointment to implement inter-domain SAV together, we suppose they agree on and know their AS relationships with each other. However, even if they do not know, they can attain these AS relationships using above algorithms.

4. Architecture of AS Relationship Based Inter-domain Filtering (ARBIF)

4.1. Overall Architecture

This section describes the architecture of ARBIF. The ARBIF system mainly consists of two components with different functions: the Validation Router (VR) and the AS-IP Prefix Mapping Server (AIMS). Border routers in an AS act as its VRs, while AIMS is a global infrastructure working for all ASes in the system. An example of the ARBIF system is shown in Figure 2.




Figure 2: An example of the ARBIF system

We adopt the ASN SAV rules from neighboring ASes to facilitate the update, propagation, and computation of these rules with BGP UPDATE messages. We also utilize the authorized address mapping provided by the centralized server AIMS to compute the required IP prefix SAV rules.

4.2. Validation Router (VR)

4.2.1. VR Role in the ARBIF system

Existing AS border routers act as Validation Routers (VR) in the ARBIF system. VRs SHOULD actively advertise their ASN SAV rule updates to neighbors according to their AS relationships and export rules when the rules change. When receiving neighbors' ASN SAV rule updates, they SHOULD decide whether to update their ASN SAV rules accordingly. VRs SHOULD also communicate with AIMS regularly to fetch IP prefixes owned by certain ASes. After several advertisements and updates, ASN SAV rules in these VRs gradually converge. VRs translate them into IP prefix SAV rules using fetched IP prefixes. Finally, VRs filter incoming packets with IP prefix SAV rules.

Each VR records its ASN SAV rules and IP prefix SAV rules, which indicate the validity of source ASes and IP prefixes. It stores these rules in the Neighbor SAV Rule Table to implement ARBIF, because VRs use them to filter spoofed packets at the AS and prefix level.

The Neighbor SAV Rule Table in a VR also stores other related information of its neighbor ASes. Table 4 shows one specific example of the Neighbor SAV Rule Table. Specifically, the table records AS numbers, relationships, connected interfaces, corresponding ASN SAV rules, and IP prefix SAV rules.

Interface	ASN	AS Relationship	ASN SAV rules	IP prefix SAV rules
Int 1	ASN 1	P2P	ASN 4	P4, P5
Int 2	ASN 2	P2C	ASN 5	P6

Table 4: An example of Neighbor SAV Rule Table

4.3.2. AIMS Implementation based on RPKI

In current networks, we choose to use RPKI as the trust anchor in our system and use Relying Party (RP) to obtain RPKI objects. Each AS deploys RP in different ways, but they all can provide ROA objects to AS border routers. Therefore, RPKI can provide the necessary information for VRs in the ARBIF system. Figure 3 shows an example of the ARBIF system based on RPKI.




Figure 3: An example of ARBIF system based on RPKI

RP can synchronize the data in RPKI repositories to local caches at regular intervals and provide objects, such as ROAs, to border routers through the RTR protocol. According to [RFC9582], an ROA is a digitally signed object that records which AS is authorized to originate one or more particular IP address prefixes. The main contents recorded in ROA are shown in Table 6. Although one ROA object can record more than one IP prefix, IP prefixes that an AS is authorized to originate may be recorded in multiple ROA objects in many cases.

...	ASN	IP Prefix	Max Length	(IP Prefix 2)	(Max Length 2)	...
-----	-----	-----------	------------	---------------	----------------	-----

Table 6: Main Content of a ROA Object

By combining all ROAs, we can obtain a full view of the IP prefixes that each AS is authorized to originate, which is the mapping information required by our ARBIF system (as shown in Figure 4). ROAs can provide authorized relations of ASNs and IP prefixes. However, to apply them to our ARBIF system, it is necessary to query further and integrate ROA objects, which reflects the necessity of AIMS.

```

AS Number 1
| -- IP Prefix 1
| -- IP Prefix 2
| -- IP Prefix 3
AS Number 2
| -- IP Prefix 4
| -- IP Prefix 5
.....

```

Figure 4: Mapping from ASNs to IP Prefixes which ARBIF needs

To meet corresponding requirements, the ARBIF system SHOULD integrate the obtained ROA objects, generate a mapping from ASNs to IP prefix sets, and provide it to VRs. This process can be implemented in RPs. RPs regularly synchronize ROA objects from the RPKI repository, integrate them, and transfer the data to VRs for them to generate IP Prefix SAV rules. In a possible design, AIMS is implemented based on the RPKI with an additional integration function in RPs. Its schematic process is shown in Figure 5.

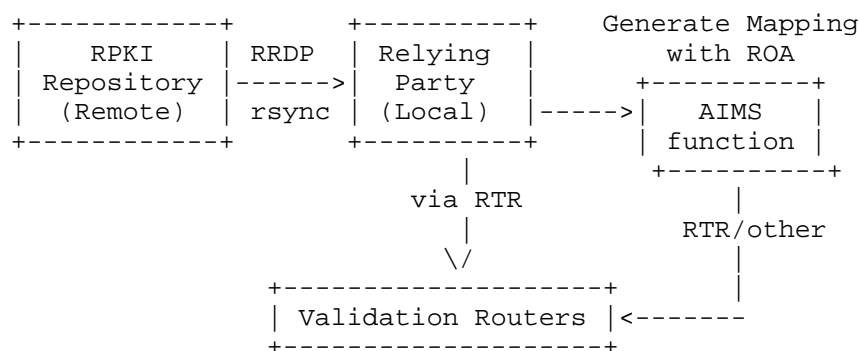


Figure 5: A schematic process of AIMS based on RPKI

4.3.3. Lightweight AIMS Implementation without RPKI

In addition to implementations based on RPKI, in some scenarios, AIMS can also be directly implemented as lightweight servers maintaining the mapping from ASNs to IP prefixes. If the traffic and connection conditions of several neighbor ASes are stable and not complex, when they deploy inter-domain SAV together but have not yet deployed RPKI, a lightweight AIMS server can be deployed first. This AIMS can maintain address mappings of these neighbor ASes, and obtain those of other related ASes using some public services.

5. BGP Extension for Inter-domain SAV

[inter-domain-sav-archt] mentions that the SAV-specific communication information mechanism can be implemented by a new protocol or an extension to an existing protocol. Following its ideas, we propose an implementation for our ARBIF scheme by an extension to BGP in this section.

5.1. Feasibility of BGP Extension

As [RFC4271] states, the primary function of a BGP speaking system is to exchange network reachability information with other BGP systems. Each router advertises route paths to networks it can reach. After further propagation, each router establishes a routing table with BGP routes to its reachable networks. The ARBIF system uses BGP forwarding routes to approximate the reverse routes. Those VRs can deploy it to calculate networks that can reach them.

Routers do not regularly announce their routing tables but incrementally advertise them using BGP UPDATE messages when they are updated. ARBIF calculates and updates SAV rules at the AS level, taking AS relationships as the abstract of inter-domain routing information. After convergence, only when increases, decreases, or changes occur to their AS relationships with neighbor VRs do they update their ASN SAV rules and advertise the updates to neighbor VRs. At this time, changes also occur to their routing tables, and they will send BGP UPDATE messages to neighbor VRs. Therefore, VRs can advertise ASN SAV rule updates with BGP UPDATE messages.

All these allow us to implement the ARBIF system with the existing route mechanism and advertise ASN SAV rule updates using BGP UPDATE messages.

5.2. Implementation of BGP Extension

To achieve the goal, we can slightly modify BGP UPDATE messages, enabling it to complete the advertisement of ASN SAV rules when advertising updated routes.

Every BGP UPDATE message contains withdrawn routes, path attributes, and Network Layer Reachability Information. The path attribute part is a sequence of BGP path attributes and can carry many attributes in one message. Each path attribute is recorded as a variable-length triple <Type, Length, Value>, allowing for various information transfers. What's more, new path attributes can be registered after IANA allocates new type codes to them.

All these above allow us to design a new BGP path attribute to exchange ASN SAV rules between AS border routers. With this attribute, an AS border router deploying ARBIF can use BGP UPDATE messages to advertise corresponding ASN SAV rule updates while updating routing information. We name it SAV_INFO for now.

SAV_INFO is a triple <attribute type, attribute length, attribute value> referring to the format proposed in [RFC4271]. The attribute type is a two-octet field containing some flags and an allocated type code. The value field records ASN SAV rules containing one or more AS numbers, each encoded as a 2-octet length field. The length field is the length of the value field in octets, occupying one or two octets.

The later section will use a concrete example to demonstrate the BGP extension for the ARBIF scheme. Our follow-up drafts will discuss the detailed implementation of this design.

5.3. An example of BGP Extension

Figure 6 shows a simple example network. After the VRs of AS 1 and AS 2 establish a BGP connection, AS 1's VR advertises its route for prefixes P1 with BGP UPDATE messages. If AS 1 deploys our ARBIF system, its VR will also announce its ASN SAV rules to AS 2's VR in these BGP UPDATE messages. AS 2's VR also advertises its information to AS 1's VR.




Figure 6: The initial example network

Figure 7 shows updates on this network. AS 3 is a new AS connected to AS 1 as a customer AS. Through its connection with AS 1, its VR advertises its routes for P3 to AS 1. AS 1's VR thus learns new routes for P3 through AS 3 and new SAV rules.




Figure 7: The updated example network

Considering BGP and SAV mechanisms, AS 1 should propagate routes and ASN SAV rules from AS 3 further to its neighbors. Since AS 3 is AS 1's customer and AS 1 is AS 2's customer, according to the export rules in Table 2, AS 1 should advertise the routes and ASN SAV rules learned from AS 3 to AS 2.

AS 1's VR propagates its newly learned routes using BGP UPDATE messages. The message's NLRI field carries the prefix P3, and the Path Attributes field adds our SAV_INFO field. Its AS_PATH attribute records the path to AS 3 through AS 1. Its SAV_INFO attribute carries the AS 3's AS number as AS 1's updated SAV rules. Receiving this BGP UPDATE message, AS 2's VR can learn the routes for AS 3 and updated ASN SAV rules.

This example shows that the ARBIF system can utilize BGP UPDATE messages to complete the ASN SAV rule advertisement while propagating the inter-domain routes.

6. Scenarios

6.1. Multi-homing Scenarios

In this section, we utilize some use cases as examples to show that our inter-domain SAV system, ARBIF, performs well in multi-homing scenarios. Our SAV scheme performs a lower false positive rate than existing mechanisms, filling the research gap proposed in [inter-domain-sav-ps].

6.1.1. Multipoint Interconnection Scenario

In other particular multi-homing scenarios, ARBIF can complete inter-domain SAV at the AS level. Figure 8 presents a scenario of multipoint interconnection between ASes. In this example, AS 1 connects with AS 2 through two pairs of VRs. AS 1 and AS 2 are in a hybrid relationship, and AS 2 is the customer of AS 2 at point 1 while they are peers at point 2.




Figure 8: An example of multipoint interconnection

If AS 1 has deployed the ARBIF system, VRs at points 1 and 2 will allow AS 2 as a source AS at the AS level. Meanwhile, according to AS 2's IP address prefixes recorded in AIMS, they will allow all these prefixes as source IP prefixes at the prefix level. At their interfaces facing AS 2, VRs at points 1 and 2 use allowed IP prefixes to filter incoming packets.

6.1.2. Multi-homing Scenario

In multi-homing scenarios, the ARBIF system improves the validation accuracy in customer interfaces, filling the gap of false positives proposed in [inter-domain-sav-ps].

We take Figure 9 as an example to analyze how ARBIF solves the limited propagation of prefixes. This figure presents a multi-homing scenario where uRPF mechanisms may lead to the problem. In this scenario, AS 2 and AS 1 are providers of AS 3, and AS 1 is the provider of AS 2. AS 3 adds the NO_EXPORT community attribute to all BGP advertisements to AS 2, preventing AS 2 from further propagating its prefixes.

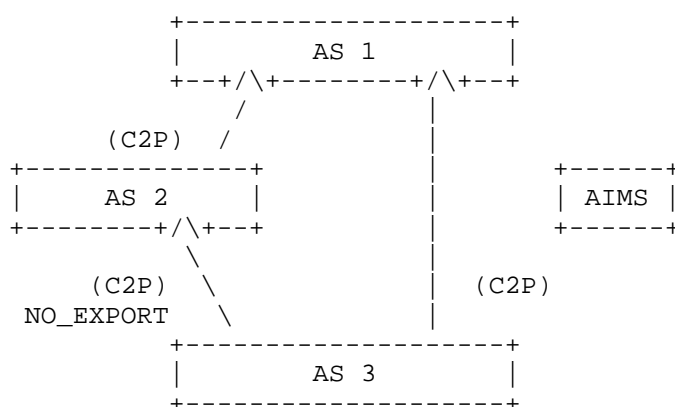


Figure 9: An example of multi-homing scenario

When deploying uRPF mechanisms, the VR facing AS 2 in AS 1 may improperly block packets originating from AS 3. If it deploys the ARBIF system, it will generate SAV rules using ASN SAV rules transmitted between VRs. When determining whether ASN SAV rules should be further propagated, BGP attributes have no effect. However, ASN SAV rule propagation depends on BGP UPDATE messages and is affected by their limitations. Since we hope that ASN SAV rules advertisement can ignore fine-grained factors, we tend to use

additional BGP UPDATE messages as a supplement to advertise SAV rules in special cases. Therefore, the VR will allow those packets originating from AS 3 to pass, avoiding false positives.

As for the problem of hidden prefixes, we solve it by specially setting the initial SAV ASN rules advertised by each AS. Under some circumstances, one AS may have particular settings and send packets with source addresses that it does not advertise, like direct server return (DSR). If deploying ARBIF, its VRs initially advertise the origin ASNs of all possible legitimate packets it can send. Therefore, these VRs will allow packets that match the specific configurations to pass, effectively avoiding false positives.

Besides false positives, [inter-domain-sav-ps] also points out false negatives within AS customer cones. The ARBIF scheme does not propose a targeted solution for this gap but does propose some ideas. A system on the data plane for traffic monitoring and management may help with limiting attacks within customer cones. What's more, in the SAVA architecture proposed in [RFC5210], access network and intra-domain SAV can prevent source address spoofing within AS and help to reduce attacks within customer cones.

6.2. Dynamic Scenario

This section utilizes some designed use cases to show how our ARBIF system performs in different dynamic scenarios. This ARBIF system handles updates at the AS level and ignores more fine-grained route updates. It reduces rule update frequency at the cost of tiny false negatives, cutting down the SAV system's update overhead.

We take the network shown in Figure 10 as an example before all changes happen. In this example, AS 1 is AS 2 and AS 3's provider, and all ASes have deployed the ARBIF system. When diverse changes occur to this network, we show the network after changes and discuss the updates of the ARBIF system.




Figure 10: An example network before changes happen

6.2.1. AS Relationships Change

Figure 11 displays the example network after AS relationships change. If the AS relationship between AS 1 and AS 2 changes from C2P to P2P, the VR in AS 1 facing AS 2 and the VR in AS 2 facing AS 1 will modify the AS relationships in their Neighbor SAV Rule Tables and remove previous SAV rules. These VRs will actively advertise their new ASN SAV rules to neighbors. VRs further propagate these rules through VR connections until they come to a new convergence in the changed network.




Figure 11: The example network after AS relationships change

6.2.2. AS Prefixes Change

Figure 12 displays the example network after AS prefixes change. If AS 3's IP address prefixes change from P3 to P4, VRs will modify the SAV information about AS 3 in their Neighbor SAV Rule Tables. Under this circumstance, VRs' ASN SAV rules remain unchanged, but they will adjust IP prefix SAV rules according to the new mapping recorded in AIMS.

In our ARBIF system, VRs use ASN SAV rules as advertised SAV rules. VRs translate ASN SAV rules into IP prefix SAV rules with the mapping provided by AIMS and do not further propagate prefix ones. Therefore, AS prefixes change won't break achieved convergence. In this example, the change of AS 3's prefixes does cause VRs to update their SAV information about AS 3. However, all ASN SAV rules remain unchanged, and VRs only update IP prefix SAV rules about AS 3.




Figure 12: The example network after AS prefixes change

6.2.3. AS Network Topologies Change

Figure 13 displays the example network after the AS network topology changes. If the AS connections change and AS 3 becomes AS 2's peer from AS 1's customer, AS 2 will add one new VR, and AS 3 will adjust its original VR. After reconfigurations, added VR in AS 2 and adjusted VR in AS 3 will fill in their Neighbor SAV Rule Tables according to the latest network situation. These VRs will actively advertise their new ASN SAV rules to neighbor ASes. VRs further propagate ASN SAV rules through VR connections until they come to a new convergence in the changed network.




Figure 13: The example network after AS network topologies change

6.2.4. BGP Attributes Change

Figure 14 displays the example network after BGP attributes change. If the BGP attributes between AS 1 and AS 2 change while other information does not, all VRs in the network needn't update their SAV information. In this example, AS 2 adds the NO_EXPORT community attribute to all BGP advertisements from it to AS 1, preventing AS 1 from further propagating its prefixes. Routing information propagated from AS 1 to AS 3 changes and no longer contains routes to AS 2.

However, our ARBIF system does not consider BGP attributes when determining whether to further propagate ASN SAV rules. In this case, when route updates are propagated with BGP UPDATE messages, ASN SAV rules will not be modified. Therefore, AS 3's ASN and IP prefix SAV rules remain unchanged, as do other ASes'.

The results indicate that our SAV scheme ignores fine-grained routing information changes because it handles AS connections rather than BGP routes. As such processing neglects restrictions on BGP route advertisement, it may cause some additional improper permits but not additional improper blocks, which meets SAV requirements. Such processing also improves the ARBIF system's stability and lessens its update overhead.




Figure 14: The example network after BGP attributes change

6.3. IXP Scenario

IXP, which is Internet eXchange Point, is a Layer 2 LAN in the OSI network model. IXPs are built with one or many Ethernet switches interconnected together across one or more physical buildings.[learn-ixp] Some ASes pay for traffic transit by other networks. Sometimes ASes may choose to connect via IXPs to reduce costs and latency. In actual networks, IXPs are critical and distinct infrastructures. Therefore, following the comments we received, we investigated how ARBIF handles typical IXP scenarios.

We first discuss how traffic is routed in common IXPs. Many IXPs use Route Servers to reduce the number of needed BGP sessions in the IXP. RS-Clients use eBGP to advertise network reachability information to the RS, which then forwards the information to its connected RS-Clients based on its configuration, establishing routes among all its connected RS-Clients. The RS exclusively uses BGP to exchange reachability information with its RS-Clients without forwarding traffic itself. Its function is similar to a Route Reflector in an intra-domain iBGP scenario. ASes connected to one IXP generally adopt free peering when exchanging traffic. Therefore, regardless of whether one IXP is implemented using RS, the IXP scenario is essentially some peering ASes connected to each other.

With this clarified, our ARBIF just treats typical IXP scenarios as several fully-connected peering ASes. We take a simple connection scenario of an IXP in Figure 15 as an example.




Figure 15: An example of IXP scenario

AS 1, AS 2, and AS 3 in Figure 15 are three ASes connected to IXP1. When calculating SAV rules for AS 1 in the ARBIF, we treat AS 2 and AS 3 as two neighboring peers of AS 1.

7. Experiment of ARBIF Implementation

7.1. Environment

We conducted simulation experiments using the GNS3 software with the following software versions shown in Figure 16. We used a GNS3 Ubuntu image in GNS3 Docker as simulated hardware devices. In this image, we installed the open-source BIRD 2 and Routinator for our implementations.

```

VMware-workstation: 17.5.2
GNS3/GNS3 VM: 2.2.54
Ubuntu: 24.04.2 LTS
Routinator: 0.14.2
BIRD: 2.14

```


Figure 16: Experiment Environment

7.2. Implementation Method

Based on the test environment, we performed experiments following the implementations described in Section 4. A simple topology, as shown in Figure 17, is used in our simulation experiments. First, we use BIRD as the implementation base of VR and do some configurations and extensions to implement ARBIF on this basis. At the same time, according to the descriptions of VR in Section 4.2.2, we used BGP Roles based on BIRD to enable VRs to obtain AS relationships between neighbor ASes.




Figure 17: A Simple Network used in the Experiment

Furthermore, we selected Routinator, a third-party RP software written in Rust, as the RP, and RTRlib with BIRD as the implementation for receiving RTR packets in VRs. Routinator regularly synchronizes necessary objects from RPKI repositories to local caches, integrates them, and sends them to VRs through the RTR protocol. Therefore, VRs can obtain ROA objects through the process mentioned in Section 4.3.2. Through a similar process, it is also feasible for VRs to obtain ASPA objects.

Judging from current implementations, we have avoided high update costs brought by new devices or new protocols. Instead, we extended existing mechanisms, namely BGP Roles, RPKI, and the BGP protocol, as the ARBIF system implementation.

7.3. Simulation Experiment Result

We have completed the ARBIF implementation in these simulation experiments. In simple simulation scenarios, the ARBIF implementation has achieved our goal in inter-domain source address validation.

8. Considerations on Deployability

8.1. Utilize existing information as much as possible

Using information beyond existing will inevitably incur additional costs due to its need for upgrades. At the same time, it will improve the deployment requirements, which prevent SAV schemes' large-scale promotion. Therefore, an easily deployable SAV scheme in real networks always utilizes existing information as much as possible. Similarly, when existing facilities can provide certain services, deployable solutions always prefer to use them rather than establish new ones.

For SAV schemes, existing information includes routing information, business relationships between different ASes, and the mapping from ASNs to IP address prefixes provided. Existing facilities include RPKI and AS border routers. The ARBIF scheme establishes the SAV system with the existing information and devices.

8.2. Prefer to use and exchange more abstract information

Unlike fine-grained concrete information, abstract information lacks details but fundamentally simplifies problems. However, it can reduce computational costs and improve efficiency, which is more conducive to promoting SAV deployment. When multiple SAV nodes collaborate, they can exchange abstract rules and generate fine-grained ones when setting prefix filters.

As discussed above, AS relationships determine the routing information between ASes and are more abstract than that. Therefore, our inter-domain SAV scheme uses AS relationships instead of routing information to calculate SAV rules at the AS level. It transmits ASN SAV rules between ASes instead of IP prefix SAV rules and only generates IP prefix SAV rules in VRs using ASN SAV rules.

8.3. Balance accuracy, time and cost

When the network remains stable, directly generating the most accurate filtering rules during forwarding table establishment is the best idea. However, the Internet often changes at different levels, which triggers validation rule fluctuations until they reconverge. We have discussed some changes and their impacts in Section 6.2. Long convergence time is not conducive to providing validation support in a constantly changing network. Therefore, an easily deployable validation scheme in the dynamic network should balance convergence time and accuracy.

When rule calculation and deployment do not bring additional costs, using the most accurate algorithms is the most effective. However, SAV schemes that need more data and calculations often have higher costs in real networks. Trading excessive expenses for a slight accuracy improvement is an inappropriate choice. Therefore, an easily deployable SAV scheme in practical situations should balance computational cost and accuracy.

The above analyses of two examples indicate that different evaluation metrics may have hidden contradictions in practical networks, making it difficult to optimize them simultaneously. The ARBIF scheme tries to balance accuracy, time, and cost.

9. Next Step

The current discussion and design do not cover all details. For example, we discuss the major and complex AS relationships in Section 3, but do not consider other complex and minor ones. In future research, we hope to obtain more complex AS relationships and connection scenarios. We will apply current system design and implementations to more AS relationships and practical scenarios. By analyzing the results, we can further optimize our ARBIF system and supplement it for special cases. We will also further refine our ARBIF implementations, enhance their security and efficiency, and reduce their overhead.

10. Security Considerations

The security considerations of our ARBIF scheme are similar to those of [inter-domain-sav-archt].

11. IANA Considerations

This draft proposes using a new BGP attribute to carry ASN SAV rules. The new BGP attribute needs an attribute type code assigned by IANA. We will put forward specific IANA considerations in a further draft about the BGP attribute implementation.

12. References

12.1. Normative References

- [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>>.

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC9234] Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages", RFC 9234, DOI 10.17487/RFC9234, May 2022, <<https://www.rfc-editor.org/info/rfc9234>>.
- [RFC9582] Snijders, J., Maddison, B., Lepinski, M., Kong, D., and S. Kent, "A Profile for Route Origin Authorizations (ROAs)", RFC 9582, DOI 10.17487/RFC9582, May 2024, <<https://www.rfc-editor.org/info/rfc9582>>.

12.2. Informative References

- [RFC5210] Wu, J., Bi, J., Li, X., Ren, G., Xu, K., Williams, M., and RFC Editor, "A Source Address Validation Architecture (SAVA) Testbed and Deployment Experience", DOI 10.17487/rfc5210, June 2008, <<http://dx.doi.org/10.17487/rfc5210>>.

`[inter-domain-sav-ps]`

Li, D., Wu, J., Liu, L., Huang, M., and K. Sriram, "Source Address Validation in Inter-domain Networks Gap Analysis, Problem Statement, and Requirements", Work in Progress, Internet-Draft, draft-ietf-savnet-inter-domain-problem-statement-08, 15 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-problem-statement-08>>.

`[inter-domain-sav-archt]`

Li, D., Chen, L., Geng, N., Liu, L., and L. Qin, "Inter-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-ietf-savnet-inter-domain-architecture-01, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-architecture-01>>.

`[sidrops-bar-sav]`

Sriram, K., Lubashev, I., and D. Montgomery, "Source Address Validation Using BGP UPDATES, ASPA, and ROA (BAR-SAV)", Work in Progress, Internet-Draft, draft-ietf-sidrops-bar-sav-06, 15 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-bar-sav-06>>.

`[sidrops-aspa]`

Azimov, A., Uskov, E., Bush, R., Snijders, J., Housley, R., and B. Maddison, "A Profile for Autonomous System Provider Authorization", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-profile-19, 6 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-profile-19>>.

`[sidrops-asra]`

Geng, N., Sriram, K., and M. Huang, "A Profile for Autonomous System Relationship Authorization (ASRA)", Work in Progress, Internet-Draft, draft-geng-sidrops-asra-profile-01, 15 April 2025, <<https://datatracker.ietf.org/doc/html/draft-geng-sidrops-asra-profile-01>>.

`[infer-relatsh]`

Gao, L., "On inferring autonomous system relationships in the Internet", December 2001, <<https://ieeexplore.ieee.org/document/974527>>.

- [as-rank] Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., and K. claffy, "AS relationships, customer cones, and validation", October 2013, <<https://dl.acm.org/doi/10.1145/2504730.2504735>>.
- [problink] Jin, Y., Scott, C., Dhamdhere, A., Giotsas, V., Krishnamurthy, A., and S. Shenker, "Stable and Practical AS Relationship Inference with ProbLink", February 2019, <<https://www.usenix.org/system/files/nsdi19-jin.pdf>>.
- [ml-pred] Varghese, J. S. and L. Ruan, "A machine learning approach to edge type inference in Internet AS graphs", April 2016, <<https://ieeexplore.ieee.org/document/7562048>>.
- [bgp2vec] Shapira, T. and Y. Shavitt, "Unveiling the Type of Relationship Between Autonomous Systems Using Deep Learning", June 2020, <<https://ieeexplore.ieee.org/document/9110358>>.
- [charact-inet] Subramanian, L., Agarwal, S., Rexford, J., and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points", June 2002, <<https://ieeexplore.ieee.org/document/1019307>>.
- [inferring-complex] Giotsas, V., Luckie, M., Huffaker, B., and K. claffy, "Inferring complex AS relationships", November 2014, <<https://dl.acm.org/doi/10.1145/2663716.2663743>>.
- [learn-ixp] Cloudflare, "What is an Internet exchange point? | How do IXPs work?", <<https://www.cloudflare.com/learning/cdn/glossary/internet-exchange-point-ixp/>>.

Acknowledgements

Thanks to Aijun Wang for his valuable comments and suggestions on this draft.

Authors' Addresses

Gang Ren
Tsinghua University
Beijing
China
Email: rengang@cernet.edu.cn

Shuqi Liu
Tsinghua University
Beijing
China
Email: liu-sq23@mails.tsinghua.edu.cn

Xia Yin
Tsinghua University
Beijing
China
Email: yxia@tsinghua.edu.cn

Minglin Jia
Tsinghua University
Beijing
China
Phone: +86 18800137573
Email: jml20@mails.tsinghua.edu.cn, millionvoid@gmail.com