SAVNET Working Group
Internet-Draft
Intended status: Informational
Expires: September 03, 2025

W. Cheng
China Mobile
D. Li
Tsinghua University
C. Lin
New H3C Technologies
S. Yue
China Mobile
X. Song
ZTE Corporation

March 03, 2025

Intra-domain SAVNET Support via IGP
draft-cheng-savnet-intra-domain-sav-igp-04

Status of this Memo

Copyright Notice

Abstract

This document introduces a new method for generating SAV rules based on the SAVNET mechanism. This method generates SAV rules layer by layer through the topology of the link state database formed by the IGP protocol.

Table of Contents

1. Introduction

In communication networks, network devices typically forward packets based only on their destination addresses, without verifying the authenticity of the source addresses. As a result, forging the source addresses raises a large number of network security problems. The main problems are as follows, as shown in Figure 1:

* Attackers attack important websites, causing them to be inaccessible and interfering with the normal use of important services by legitimate user.

* Attackers conceal their true identity and location, making it difficult to trace the source of illegal network activities.

* Attackers interfere with the normal operation of services such as accounting, management, and security authentication based on real source addresses, causing a large amount of network resources to be misappropriated.

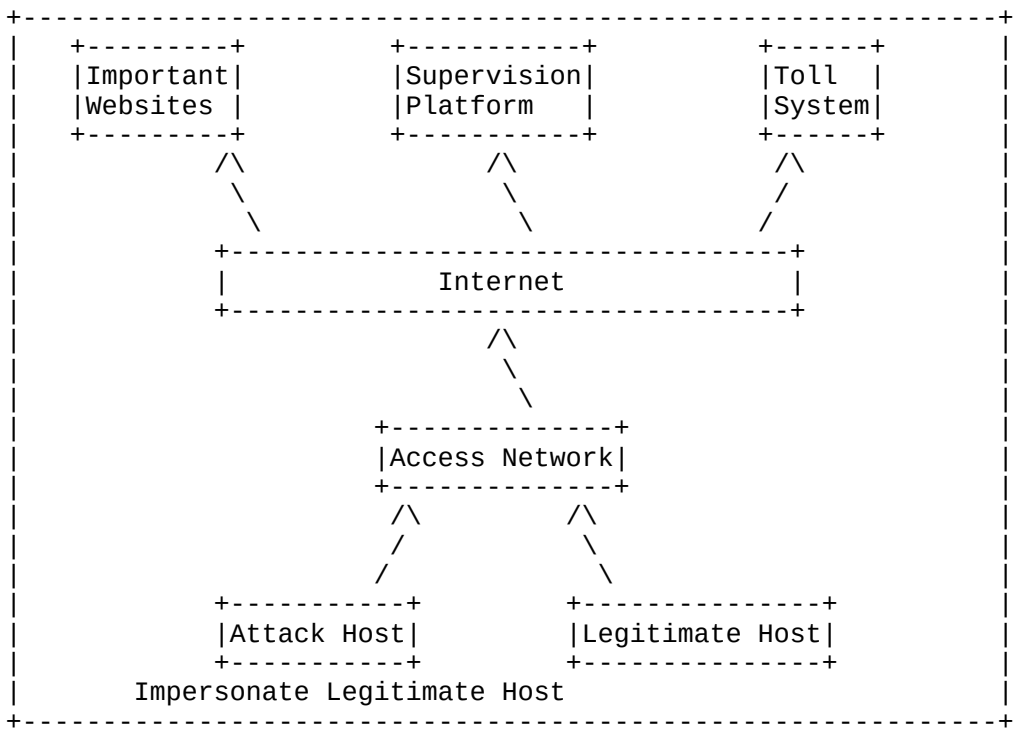

Figure 1: Scenario of source address spoofing.

Before SAVNET was proposed [I-D.ietf-savnet-intra-domain-problem-statement], several Source Address Validation (SAV) technical schemes have been proposed, such as ACL, uRPF, etc., which are dedicated to solving the illegal attacks based on source address spoofing. However, these SAV technologies still have limitations, which restrict the application of SAV technology in existing networks.

ACL SAV: This scheme can be used for both outbound traffic verification and inbound traffic verification. The ACL rules need to be updated manually in time to make them consistent with the latest filter conditions [RFC2827] [RFC3704].

Strict uRPF SAV: This scheme is typically used for outbound traffic verification. The SAV rules can be automatically generated and updated, but there is a serious problem of inappropriate blocking in asymmetric routing scenarios [RFC3704].

Loose uRPF SAV This scheme is typically used for inbound traffic verification. The SAV rules can be automatically generated and updated, but most spoofed data will be inappropriately allowed to forward [RFC3704].

In order to optimize the limitations of the above schemes, the SAVNET mechanisms based on SAV-related information are proposed. The SAVNET mechanisms, working in an incremental or partial deployment manner, can automatically adapt to network dynamics such as routing changes or prefix changes, instead of purely relying on manual update. The SAVNET mechanisms also improve the verification accuracy upon existing intra-domain SAVNET mechanisms, and allow for rapid updates of SAV rules so as to minimize the impact of improper blocking and permitting during the convergence process [I-D.ietf-savnet-intra-domain-problem-statement] [I-D.ietf-savnet-intra-domain-architecture].

This document introduces a novel method for generating SAV rules, leveraging the SAVNET mechanism. The proposed approach systematically generates SAV rules layer by layer, utilizing the topology derived from the link-state database established by the IGP protocol. This method is particularly well-suited to address the specific networking requirements of service providers.

2. Terminology

The following terminologies are used in this document.

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

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

IGP: Interior Gateway Protocol.

IGP LSDB: IGP Link-State Database.

IGP node: It is anchored by a Router-ID that is used by the underlying IGP, i.e., a 48-bit ISO System-ID for IS-IS and a 32-bit Router-ID for OSPFv2 and OSPFv3.

IGP link Each link is anchored by a pair of Router-IDs that are used by the underlying IGP, i.e., a 48-bit ISO System-ID for IS-IS and a 32-bit Router-ID for OSPFv2 and OSPFv3.

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

BFS: Breadth-First Search is a graph search algorithm. It starts at the source node and explores all the neighbor nodes at the current depth before moving on to nodes at the next depth level. BFS uses a queue to assist in the search.

3. Design Goals

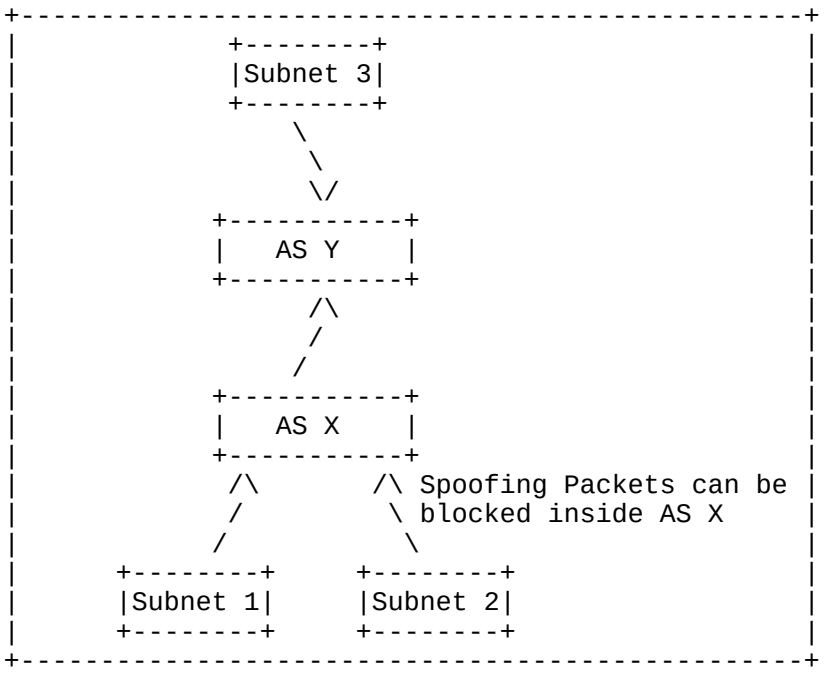

Figure 2: The case of outbound traffic verification

This method is designed to enhance the intra-domain SAVNET and achieve the following goals:

* Outbound traffic verification. As shown in Figure 2, Subnet 2 of AS X sends packets which spoof the source addresses of Subnet 1 or Subnet 3. If AS X deploys the intra-domain SAVNET solution, the spoofing packet from Subnet 2 can be blocked inside AS X.

* Incoming traffic verification. As shown in Figure 3, AS X would receive incoming traffic packets which spoofs source addresses of AS X. If AS X also can deploy intra-domain SAVNET solution, the spoofing packets from AS Y could be blocked by AS X.

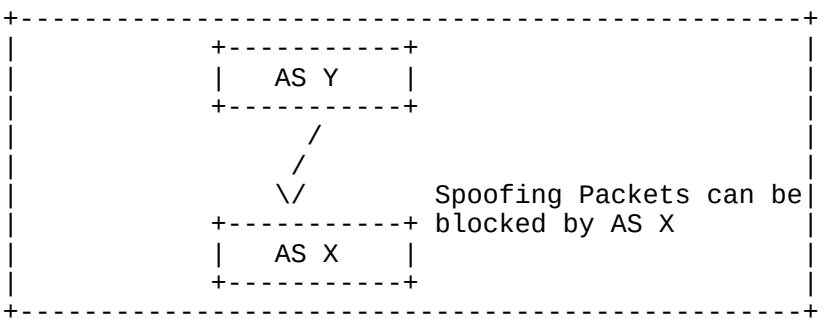

Figure 3: The case of incoming traffic verification

4. Solution

4.1. Overview

As shown in Figure 4, [draft-ietf-savnet-intra-domain-architecture-00] describes SAV filtering on three types of interfaces: interfaces on host-facing routers, such as interface 1 on Router C as shown in the figure; interfaces on customer-facing routers, such as interface 1 on Router A and interface 1 on Router B as shown in the figure; and interfaces on AS border routers, such as interface 2 on Router D and interface 2 on Router E when facing another AS.
For the first two types of interfaces, SAV filtering is primarily carried out using an allow-list. For the third type of interface, filtering is primarily implemented using a block-list.
This document also describes SAV filtering on the interfaces of routers within the domain, such as Interface 1 of Router D and Interface 1 of Router E, as depicted in the diagram. Such interfaces employ allow-list based SAV filtering, with the computation process outlined in Section 4.3.
The specific details about allow-list and block-list are detailed below.

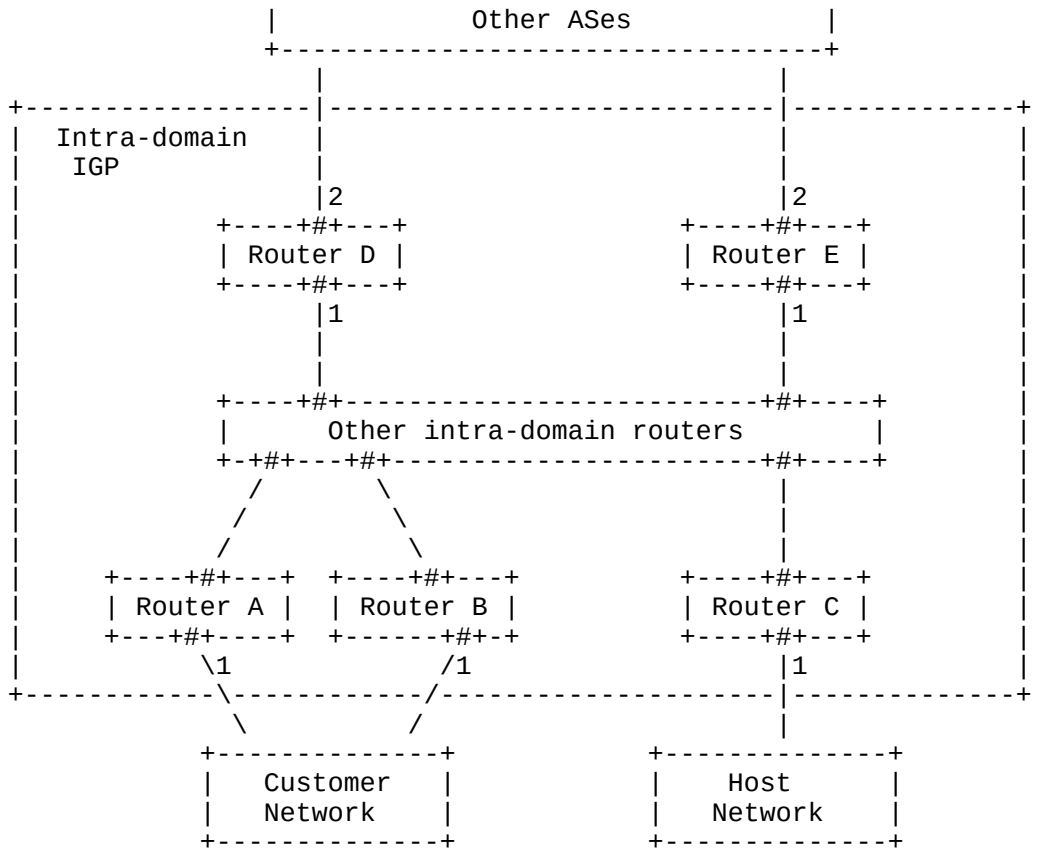

Figure 4: Overview of intra-domain SAVNET architecture

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

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

For the host-facing router interfaces, the allow-list can be generated for directly connected host segments, as shown in Figure

4. If the interface 1 network segment address of Router C is P, the allow-list (P, C-1) can be generated for Router C’s interface 1.

For the ports of routers facing the customers, as shown in Figure 4, if the interface 1 network segment address of Router A is P1, and the interface 1 network segment address of Router B is P2, the allow-lists (P1, A-1) and (P2, A-1) can be generated for Router A’s interface 1, and the allow-lists (P1, B-1) and (P2, B-1) can be generated for Router B’s interface 1.

For the router interfaces facing the customers, as shown in Figure 4, where Router A’s interface 1 belongs to network segment address P1 and Router B’s interface 1 belongs to network segment address P2, the allow-lists (P1, A-1) and (P2, A-1) can be generated for Router A’s interface 1, and the allow-lists (P1, B-1) and (P2, B-1) can be generated for Router B’s interface 1.

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

For the interface facing another AS, as shown in Figure 4, such as Router D’s interface 2 and Router E’s interface 2, this scenario typically involves using block-lists for SAV filtering. The generation of block-lists on the edge interfaces connecting to external domains can be achieved by generating block-list entries on the upper interface when learning routes from the lower interface. For example, if Router D learns the prefix P on interface 1, then a block-list entry (P, D-2) can be generated on interface 2.

4.2. Intra-domain SAVNET

This section introduces a new approach for generating SAV rules within intra-domain scenarios using the SAVNET mechanism.

The method relies on two essential pieces of information: source prefix information and reachability information.

The source prefix information indicates the origin of the source address. There are two options available: one is to treat all IGP protocol-advertised prefixes or specify certain IGP-advertised prefixes through policy as SAVNET source prefixes, while the other is to designate specific advertised prefixes as SAVNET source

prefixes by extending the IGP protocol. The extension of the IGP protocol is beyond the scope of this document.

Leveraging the acquired source prefix and reachability information, the method dynamically calculates the inbound interface information for the source addresses within the domain and generates the corresponding Source Address Validation (SAV) Rule. The general process framework of the method is illustrated in Figure 5, where the SAVNET Agent acts as the processing unit for generating SAV rules.

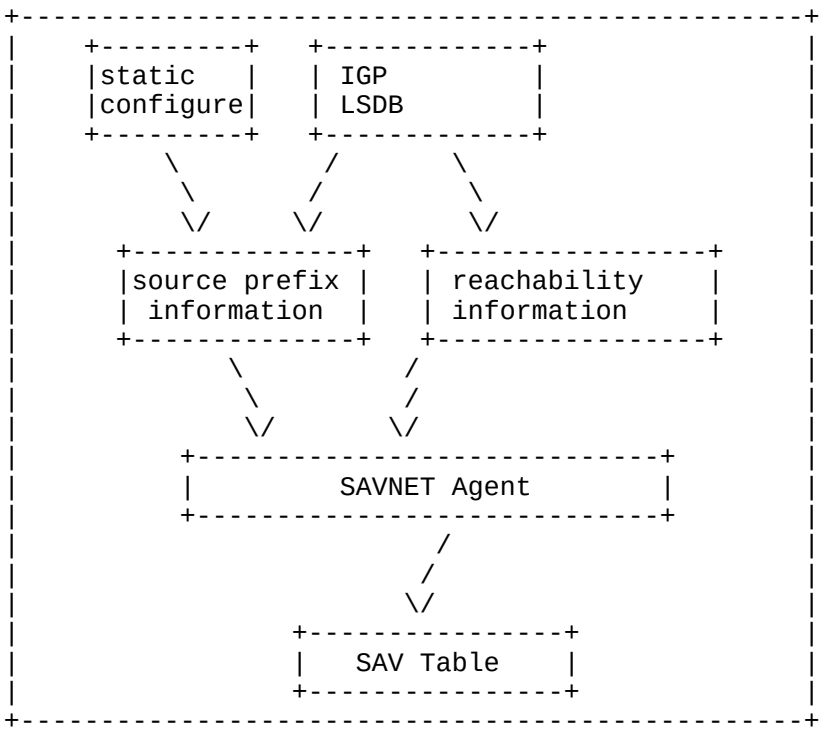

Figure 5: The overall process framework

SAVNET Agent first selects the checking node, which can be triggered by specific nodes actively requesting it or in response to changes in the topology, thereby re-checking the corresponding affected nodes.

The SAV checking node extracts connection information between nodes from the IGP Link State Database (LSDB). It traverses each link of the SAV checking node and, using topology information, determines all the node information that can be reached via that link. Finally,

the source prefix address ranges of the nodes are obtained to generate SAV Rule entries for that link. During the traversal process, topology information should not loop back to the SAV checking node.

As depicted in Figure 6, when considering the topology information of Node A, it includes the link connection L1 between A and B, as well as the link connection L2 between A and C. Upon calculating this information on Node A, the result indicates that LINK L1 can reach Node B, while LINK L2 connects to Node C.

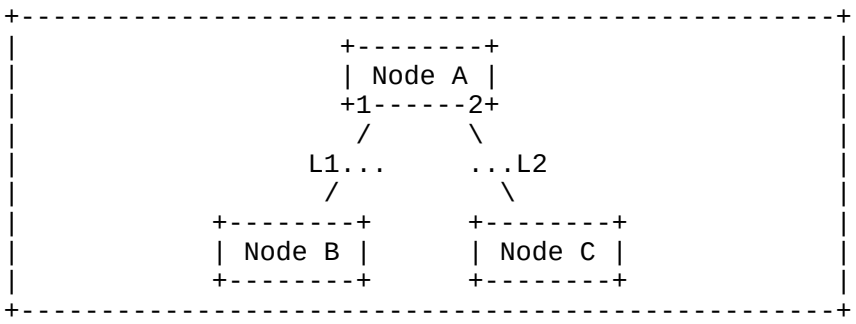

Figure 6: Example 1 of topology calculation

This method can solve the problem of intra-domain outbound and inbound traffic mentioned in Section 3, and it can automatically adjust SAV table entries according to topology changes to achieve security protection of intra-domain source addresses.

4.3. SAV RULE

4.3.1. Composition of the SAV Rule

The composition of the SAVNET rule is illustrated in Figure 7, with each SAVNET rule comprising a source prefix and a list of IGP interfaces.

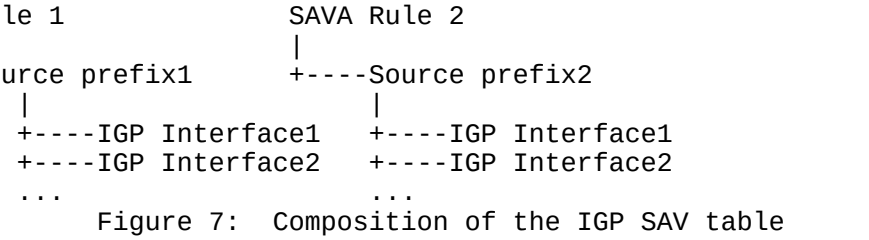

Figure 7: Composition of the IGP SAV table

4.3.2. Origin of the Source Prefix

The source prefix can be extracted from the LSDB advertised by each IGP node. This means that the source prefix information can be generated from the routing information advertised by each IGP node.

1) IS-IS Protocol

For the IS-IS protocol, IPv4 source prefixes are advertised using the "IP Extended Reach TLV", while IPv6 source prefixes are advertised using the "IPv6 Reachability TLV", without specific distinction between different types of source prefixes.

The source prefixes are categorized into Level-1 prefixes and Level-2 prefixes. When a Level-1/Level-2 node learns a Level-1 or Level-2 source prefix, it redistributes this prefix into Level-2 or Level-1 through route leaking

When computing source prefixes, only the source prefixes advertised by the current node need to be taken into account.

2) OSPF Protocol

The OSPF source prefixes can be further categorized into intra-area source prefixes, inter-area source prefixes, and external source prefixes.

* intra-area source prefix:

The intra-area source prefix can be extracted from the Router LSA, Network LSA, or OSPFv2 Extended Prefix Opaque LSA.

When computing intra-area source prefixes, only the intra-area source prefixes advertised by the current node need to be taken into account.

* Inter-area Source Prefix:
Inter-area routes are advertised by ABRs. The inter-area source prefix can be extracted from the Summary LSA or OSPFv2 Extended Prefix Opaque LSA.

When computing inter-area source prefixes, only the inter-area source prefixes advertised by the current node need to be taken into account.

* External Source Prefix:
External routes are advertised by ASBRs. The external source prefix

can be extracted from the AS-External LSA or OSPFv2 Extended Prefix Opaque LSA.

When computing External source prefixes, only the external source prefixes advertised by the current node need to be taken into account.

3) OSPFv3 Protocol

Similar to the OSPF protocol, OSPFv3 source prefixes are also categorized into intra-area source prefixes, inter-area source prefixes, and external source prefixes.

* Intra-area Source Prefix:
The intra-area source prefix can be extracted from the Intra-Area-Prefix-LSA or E-Intra-Area-Prefix-LSA.

* Inter-area Source Prefix:
The inter-area source prefix can be extracted from the Inter-Area-Prefix-LSA or E-Inter-Area-Prefix-LSA.

* External Source Prefix:
The external source prefix can be extracted from the AS-External-LSA or E-AS-External-LSA.

4.4. Procedure

The SAVNET Agent is responsible for generating SAV rules based on the source prefix and topology information. The source prefix information can be dynamically disseminated by nodes that support SAVNET functionality or manually configured on inspection nodes.

The principle for calculating topology information is as follows:

Starting from the designated start node, the SAV rules for each interface are calculated sequentially. For a specific interface, the calculation begins from that interface and traverses to find out all reachable nodes. The SAV rules for this interface, (Prefix, IF), are then generated based on the source prefixes advertised by these reachable nodes. Finally, the interface SAV rules are merged based on prefixes. The entries with the same prefix are merged into one entry indexed by the prefix with a list of interfaces. This process ensures the comprehensive generation of SAV rules based on source prefixes and reachability information, allowing for effective security enforcement within the domain.

Here is the refined and optimized process of computation based on the Breadth-First Search (BFS) algorithm:

Step 1: Before initiating the SAV rule calculation, save the existing SAV rule table to facilitate the identification of changes in SAV rule table entries.

Step 2: Traverse all interfaces of the starting node and perform SAV rule calculation for each interface. Choose an interface as the calculating interface in a sequential manner, following steps 3 to 8; calculate all the reachable nodes through this interface. Based on the source prefixes advertised by each reachable node, compute the SAV rule (Prefix, IF) for this interface.

Step 3: Clear the visited flag for all nodes and mark the starting node as visited to initialize the BFS traversal.

Step 4: Add the neighboring nodes of the calculated interface to the queue if the nodes are not visited, and mark them as visited. These neighbor nodes must be in a bidirectional connected state.

Step 5: Retrieve the first node from the queue.

Step 6: Process current node, add all adjacent unvisited nodes to the queue, and mark them as visited.

Step 7: Generate SAV rules for the calculated interface based on the source prefixes of current node. The SAV rules are indexed by prefix. Firstly, process the source prefix: If the source prefix does not exist in the SAV rules, add a new SAV rule. Then, process the interface under the source prefix: If the interface is not in the interface list, add the new interface to the list. Refer to section 4.3 for detailed information on obtaining the source prefixes for different types of prefixes.

Step 8: Repeat steps 5 to 7 until the queue is empty.

Step 9: Repeat steps 2 to 8 until SAV rules for each interface are individually calculated.

Step 10: Merge the SAV rule entries obtained by all interfaces, combining entries with the same prefix into a single entry and consolidating the interfaces from each entry into the interface list of the merged entry.

Step 11: Deploy SAV rules. By comparing the newly generated SAV rules with the saved old SAV rules, perform add/modify/delete operations on SAV rule entries.

An example of the above process is shown in Figure 8 below.

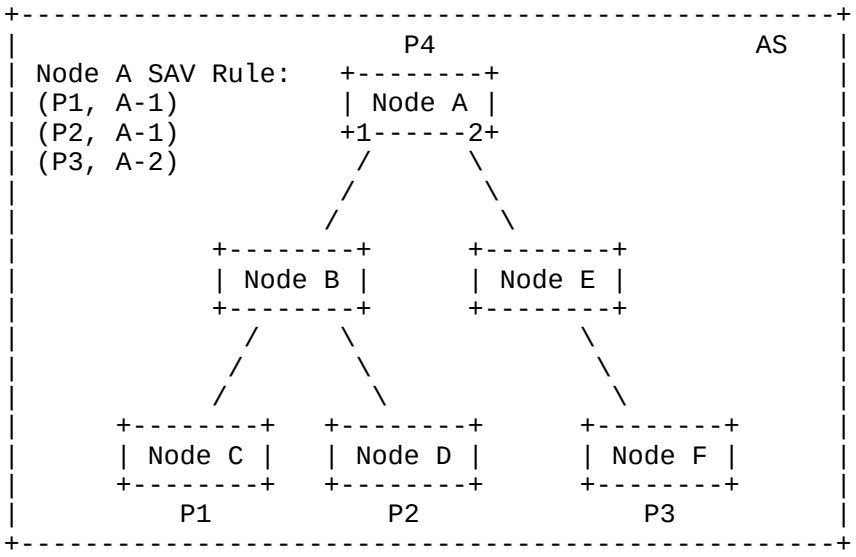

Figure 8: The case of the SAVNET procedure

The steps provided illustrate the procedure for Node A to obtain SAV rules based on the source prefix and reachability information obtained through the IGP protocol extension. The process involves traversing the network topology to gather source prefix information and generate SAV rules for each interface separately.

Based on the information presented, the following SAV rules are generated for Node A:

1) Before initiating the SAV rule calculation, save the existing SAV rule table of Node A to facilitate the identification of changes in SAV rule table entries

2) Traverse all the IGP links of Node A and generate SAV rules for each interface separately. Select interface 1 first.

3) Clear the visited flag for all notes. Mark Node A as visited.

4) Add node B to the queue and mark it as visited.

5) Retrieve node B from the queue.

6) Add all adjacent nodes of node B that have not been visited to the queue and mark them as visited. Node C and Node D are added to the queue and marked as visited.

7) Process the source prefix information of the current node to generate SAV rules. For node B, there is no source prefix advertised.

8) Retrieve node C from the queue.

9) Add all adjacent nodes of node C that have not been visited to the queue and mark them as visited. No adjacent nodes were added to the queue.

10) Process the source prefix information of node C to generate SAV rules. Generate SAV rule of prefix P1 according to the source prefix advertised by Node C, and add interface A-1 to interface list.

11) Retrieve node D from the queue.

12) Add all adjacent nodes of node D that have not been visited to the queue and mark them as visited. No adjacent nodes were added to the queue.

13) Process the source prefix information of node D to generate SAV rules, Generate SAV rule of P2 according to the source prefix advertised by Node D, and add interface A-1 to interface list.

Through the above calculation process, generate SAV rules (P1, A-1) and (P2, A-1) for interface 1 of node A.

Similarly, for interface 2 of node A, the SAV rule (P3, A-2) can be generated.

4.5. Multi-Homed Scenarios

In the multi-homed scenario depicted in Figure 9, the prefix 10.0.0.0/16 is accessed via nodes D and E, with the subnet 10.0.1.0/24 being accessed via node D, and 10.0.0.0/24 being accessed via node E. Through IGP connectivity calculation, the SAV table entries derived on node A are: (10.0.1.0/24, A-1), (10.0.0.0/24, A-2).

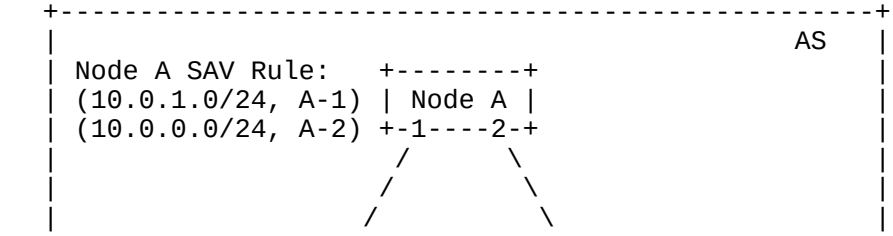

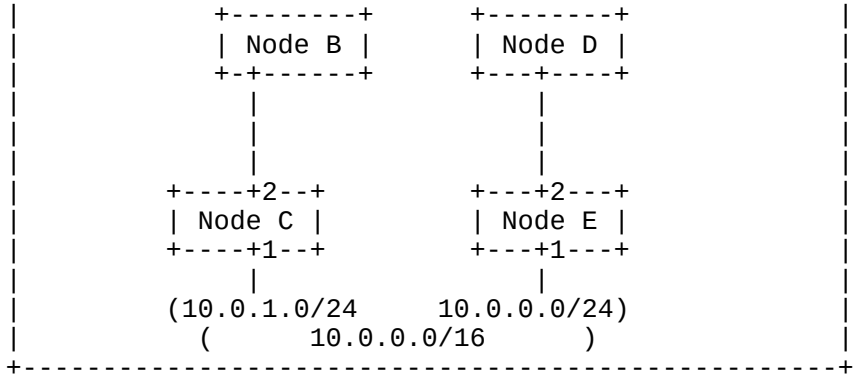

Figure 9: The case of multi-homing access

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

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

To address this issue , one approach is to configure a static route (10.0.0.0/24, C-1) on the interface C-1 of Node C, and specify a higher Cost value than the route on Node E. Similarly, on Node E, configure a static route (10.0.1.0/24, E-1) with a higher Cost value compared to the route on Node C.
The resulting SAV entries are as follows:
Entries on Node C: (10.0.1.0/24, C-1), (10.0.0.0/24, C-1)
Entries on Node E: (10.0.0.0/24, E-1), (10.1.0.0/24, E-1)
Entries on Node A: (10.0.1.0/24, A-1), (10.0.0.0/24, A-1)
(10.0.1.0/24, A-2), (10.0.0.0/24, A-2)

4.6. Peer Devices Scenarios

In the scenario depicted in Figure 10, to enhance network fault tolerance, two identical devices are deployed at equivalent positions in the network. These devices are directly connected by a cross-link, which refer to as the peer-to-peer devices, with a direct link between the two peer devices termed a cross-link.

When peer devices are present, if connectivity is calculated according to intra-domain protocols, the source prefix information calculated on interfaces 1 and 2 of Node A, as well as interfaces 1 and 2 of Node B, will be identical. This results in the loss of source prefix filtering effectiveness. To avoid this situation, during connectivity calculations, if the cross-link between the peer devices is in the UP state, connectivity calculations will treat the peer devices as a single entity. Connectivity calculations for both peer devices are initiated only when the cross-link between the peer devices is in the DOWN state.

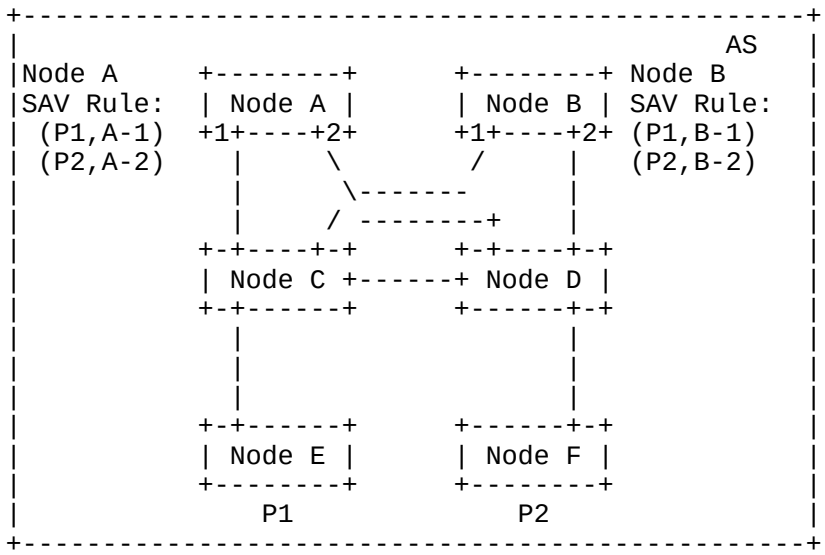

Figure 10: The case of peer devices

The corresponding processing procedure requires a modification to Step 4 in Section 4.4:

Step 4: If the neighboring nodes of the calculated interface have not been visited, add them to the queue and mark them as visited. These neighboring nodes must be in a bidirectional connected state. If these devices have corresponding peer devices, mark their peer devices as visited as well.

5. Example

The intra-domain SAVNET method commonly applies to scenarios involving intra-domain inbound traffic and inter-domain inbound traffic. Devices close to the lower layer can intercept traffic from intra-domain nodes that forge source addresses of other nodes, thereby preventing inbound traffic attacks within the domain. Similarly, devices close to upper layer can block attack packets received from inter-domain sources that forge intra-domain source addresses, in order to protect incoming traffic from external domains.

Taking the network topology depicted in Figure 12 as an example, the source address P4 belongs to router A, the source address 10.0.0.0/24 belongs to router C, the source address 20.0.0.0/24 belongs to router D, and the source address 30.0.0.0/24 belongs to router F. All these source addresses are advertised through an IGP protocol extension, and the intra-domain path is calculated via the IGP protocol. The connecting links from routers C, D, and F to A are as follows: C->B->A, D->B->A, and F->E->A, respectively.

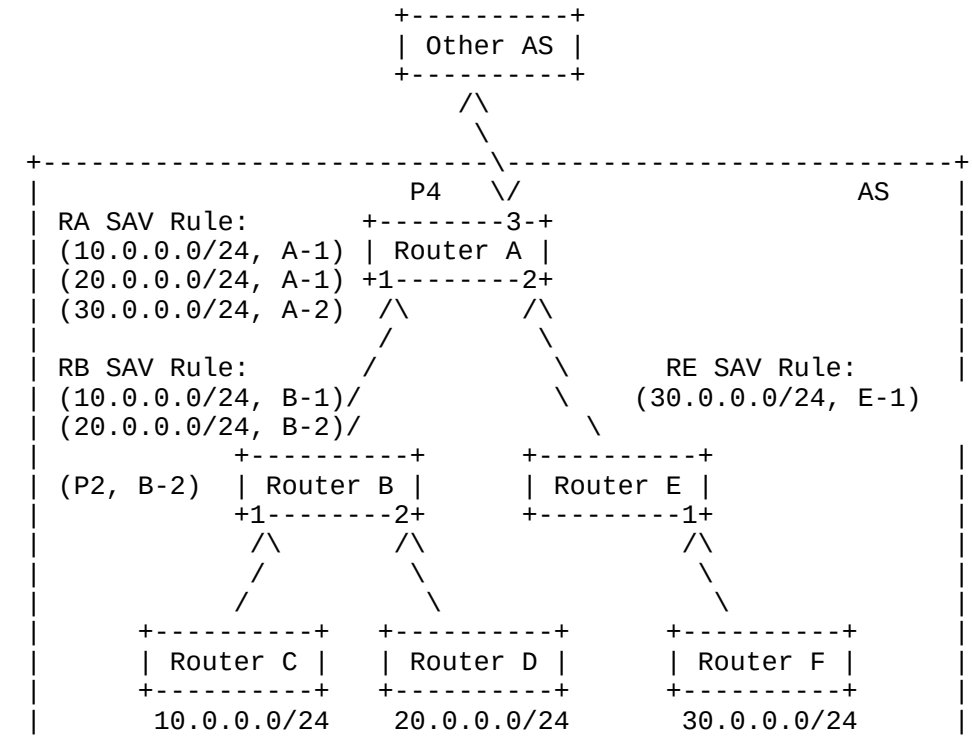

Figure 12: The intra-domain outbound traffic scenario

Based on the provided link and source address information, the intra-domain SAVNET method processes as follows:

For Router A:

The legal incoming interface for source prefixes of Router C and D is A-1.

The legal incoming interface for the source address of Router F is A-2.

For Router B:

The legal incoming interface for source prefixes of Router C is B-1.

The legal incoming interface for source prefixes of Router D is B-2.

For Router E:

The legal incoming interface for source prefixes of Router F is E-1.

Based on this information, each router in the intra-domain generates corresponding SAV rules.

With the SAVNET function enabled, the following scenarios will occur:

If Router C sends attack traffic with the source address of Router D or F, the counterfeit traffic will be intercepted by Router B.

If Router D sends attack traffic with the source address of Router C or F, the traffic with the fake source address will be intercepted by Router B.

If Router F sends attack traffic carrying the source address of Router C or D, the traffic will be blocked by Router E.

Furthermore, when Router A receives traffic from other Autonomous Systems (ASs), if the traffic forges the source addresses of intra-domain routers, Router A can intercept the traffic.

6. Manageability Considerations

The SAVNET detection feature only needs to be activated on a few selected key nodes, rather than all nodes.

Therefore, the IGP SAVNET feature is configurable. When this feature is enabled, it allows the configuration of specific source prefixes to generate SAVNET entries, which can be set through routing policies.

For Peer Devices scenarios, it allows for configuration to enable this feature.

For blacklist functionality, specific source prefixes can be configured to generate a blacklist on specific interfaces through configuration.

7. Deployment Considerations

It is desirable that all nodes in the intra-domain network could deploy this SAVNET method to automatically and accurately generate SAV rules, and therefore preventing source address spoofing attacks in the direction of outbound and inbound traffic.

However, in the existing network, only partial nodes in the intra-domain network support this method, due to asynchronous upgrades of devices. This results in that the deployed node cannot perceive the source addresses of non-deployed nodes and generate corresponding SAV rules, in spite of having all topology information. In this case of partial deployment, the deployment node can statically configure the specified source address of the non-deployment node to make up for the above shortcomings and meet the conditions for generating SAV rules.

8. IANA Considerations

TBD

9. Security Considerations

TBD

10. References

10.1. Normative References

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[I-D.draft-ietf-savnet-intra-domain-architecture]Li, D., Wu, J., Huang, M., Chen, L., Geng, N., Qin, L., and F. Gao, "Intra-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-ietf-savnet-intra-domain-architecture-00, 25 July 2023, <https://datatracker.ietf.org/doc/html/draft-draft-savnet-intra-domain-architecture-00>.

10.2. Informative References

[I-D.ietf-savnet-intra-domain-problem-statement] Li, D., Wu, J., Qin, L., Huang, M., and N. Geng, "Source Address Validation in Intra-domain Networks Gap Analysis, Problem Statement, and Requirements", Work in Progress, Internet-Draft, draft-ietf-savnet-intra-domain-problem- statement-02, 17 August 2023, https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-02>.

Acknowledgments

TBD

Authors’ Addresses

Weiqiang Cheng
China Mobile
China

Email: chengweiqiang@chinamobile.com

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

Changwang Lin
New H3C Technologies
Beijing
China

Email: linchangwang.04414@h3c.com

Shengnan Yue
China Mobile
China
yueshengnan@chinamobile.com

Xueyan Song
ZTE Corporation
China
Email: song.xueyan2@zte.com.cn