SAVNET
Internet-Draft
Intended status: Informational
Expires: 16 October 2025

D. Li
J. Wu
Tsinghua University
L. Qin
Zhongguancun Laboratory
N. Geng
Huawei
L. Chen
Zhongguancun Laboratory

14 April 2025

Intra-domain Source Address Validation (SAVNET) Architecture
draft-ietf-savnet-intra-domain-architecture-02

Abstract

This document proposes the intra-domain SAVNET architecture, which achieves accurate source address validation (SAV) in an intra-domain network by an automatic way. Compared with uRPF-like SAV mechanisms [RFC3704] that only depend on routers' local routing information, SAVNET routers generate SAV rules by using both local routing information and SAV-specific information exchanged among routers, resulting in more accurate SAV validation in asymmetric routing scenarios. Compared with ACL-based ingress filtering [RFC2827] that entirely requires manual efforts to accommodate to network dynamics, SAVNET routers learn SAV rules automatically in a distributed way.

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 16 October 2025.

Table of Contents

1. Introduction

Source address validation (SAV) is important for mitigating source address spoofing and thus contributes to the Internet security. In the Source Address Validation Architecture (SAVA) [RFC5210], SAV is divided into three checking levels, i.e., access-network SAV, intra-domain SAV, and inter-domain SAV. When an access network does not deploy SAV (such as SAVI [RFC7039][RFC7513], Cable Source Verify [cable-verify], and IP Source Guard [IPSG]), intra-domain SAV helps block spoofed packets from an access network as close to the source as possible [I-D.ietf-savnet-intra-domain-problem-statement].

The main purpose of the intra-domain SAV mechanism for an AS A, is to block the spoofing data packets from a host or customer network that use source addresses of other networks, as well as block the spoofing data packets from other ASes that use source addresses of AS A. The main task of the intra-domain SAV mechanism is to generate the correct mapping relationship between a source address (prefix) and the valid incoming router interface(s), called SAV rules. The core challenge of the intra-domain SAV mechanism is how to efficiently and accurately learn the mapping relationship. Although many existing intra-domain SAV mechanisms (such as ACL-based ingress filtering [RFC2827], strict uRPF [RFC3704], and loose uRPF [RFC3704]) have been proposed, they suffer from either inaccurate mapping in asymmetric routing scenraios, or high operational overhead in dynamic networks. The key cause is that exsiting mechanisms generate the SAV rules by a router's local routing information or by manual inputs. To addresses problems of existing intra-domain SAV mechanisms, five requirements for a new intra-domain SAVNET mechanism are proposed in [I-D.ietf-savnet-intra-domain-problem-statement].

This document introduces the intra-domain SAVNET architecture to meet the five requirements and guide development of future intra-domain SAV solutions. The key idea of intra-domain SAVNET is to generate SAV rules in routers based on SAV-specific information exchanged among routers, instead of solely depending on local routing information like in existing mechanisms. It achieves accurate SAV validation, because SAV-specific information is specialized for SAV and thus helps generate more accurate SAV rules than solely using local routing information. It achieves automatic SAV rule update, because SAV-specific information exchange is triggered when there is topology change or prefix change. In the incremental/partial deployment scenario where only part of intra-domain routers support the intra-domain SAVNET, it provides incremental benefits by using SAV-specific information provided by routers that support the intra-domain SAVNET, and/or local routing information to generate SAV rules.

The reader is encouraged to be familiar with [I-D.ietf-savnet-intra-domain-problem-statement] and [huang-savnet-sav-table].

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

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

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

SAV-related Information: The information used by a router to make SAV decisions. For intra-domain SAV, SAV-related information includes both local routing information and SAV-specific information.

SAV-specific Information Communication Mechanism: The mechanism for exchanging SAV-specific information between routers. It can be either a new protocol or an extension to an existing protocol.

SAV Information Base: A table or data structure in a router which stores SAV-specific information and local routing information.

SAV Rule: The rule in a router that describes the mapping relationship between a source address (prefix) and the valid incoming interface(s). It is used by a router to make SAV decisions.

SAVNET Router: An intra-domain router which runs intra-domain SAVNET.

SAVNET Agent: The agent in a SAVNET router that is responsible for communicating SAV-specific information, processing SAV-related information, and generating SAV rules.

Host-facing Router: An intra-domain router facing an intra-domain host network.

Customer-facing Router: An intra-domain router facing an intra-domain customer network.

AS Border Router: An intra-domain router facing an external AS.

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.

3. Overview

Figure 1 illustrates intra-domain SAVNET architecture in an intra-domain network. To generate more accurate SAV rules, intra-domain SAVNET allows SAVNET routers to automatically exchange SAV-specific information. Every SAVNET router can choose which SAVNET routers to provide its SAV-specific information to. Arrows in Figure 1 indicate the direction of SAV-specific information flows originated from Router A and Router C. SAV-specific information flows originated from other routers are omitted for brevity. After receiving SAV-specific information provided by other routers, the SAVNET router can generate more accurate SAV rules by using SAV-specific information provided by other routers, its own SAV-specific information, and/or routing information in the local FIB/RIB.

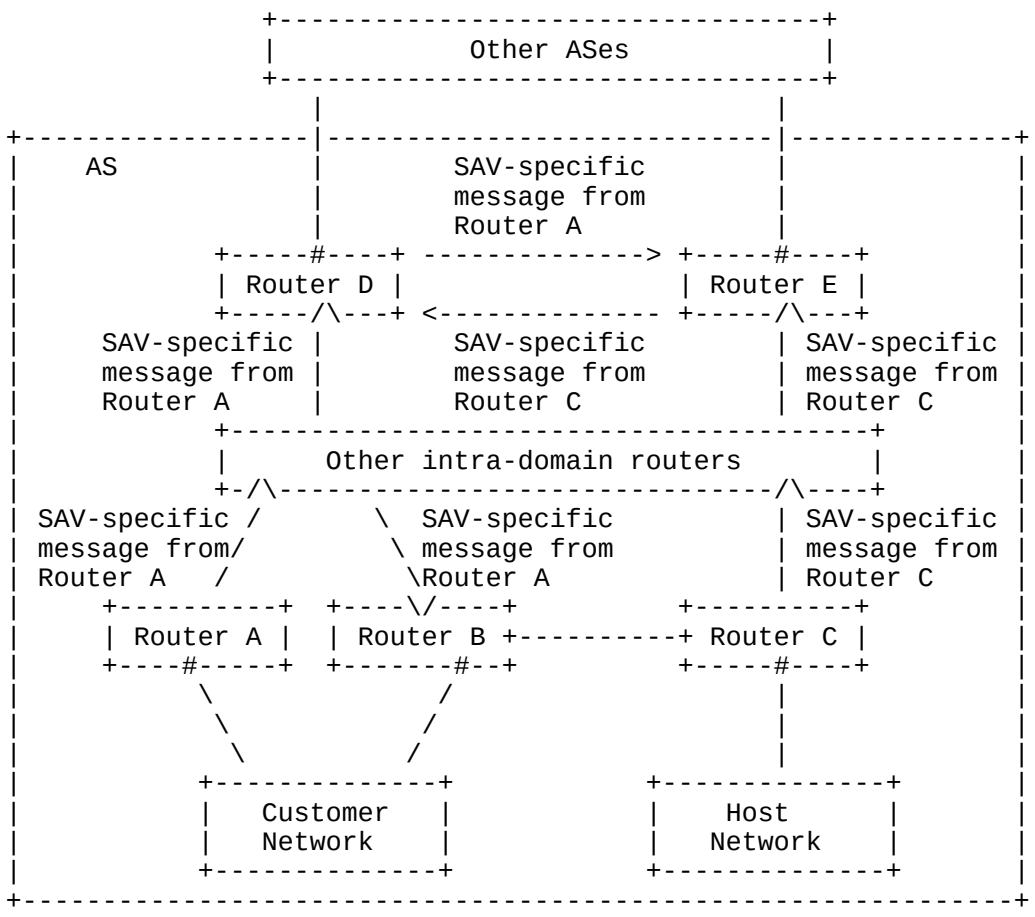

Figure 1: Overview of intra-domain SAVNET architecture

For example, the host-facing (or customer-facing) router can contain the locally-known source prefixes of the network it is facing in its SAV-specific information and provide its SAV-specific information to other routers. When Router B receives Router A's SAV-specific information, it can learn all source prefixes belonging to the customer network in combination with its locally-known source prefixes of the customer network, even if there is an asymmetric route between Router B and the customer network. After that, Router B can block source-spoofed data packets from the customer network that use source addresses not belonging to the customer network. Routers D and E can identify source prefixes belonging to the local AS by using SAV-specific inforamtion provided by Routers A, B, and C. They can block source-spoofed data packets from other ASes that use source addresses belonging to the local AS.

4. Roles of SAVNET Routers

Every SAVNET router has a SAVNET Agent that is responsible for actions related to SAV. As shown in Figure 2, a SAVNET router can act as one or two roles in the intra-domain SAVNET architecture, namely, source entity to provide its SAV-specific information to other SAVNET routers, or/and validation entity to receive SAV-specific information from other SAVNET routers.

4.1. Source Entity

When a SAVNET router acts as source entity, the information provider of its SAVNET Agent provides its SAV-specific information to other SAVNET routers that act as validation entity. For example, a host-facing router acting as source entity can obtain its SAV-specific information related to the host network to which it is connected and selectively provide this information to other SAVNET routers.

4.2. Validation Entity

When a SAVNET router acts as validation entity, the information receiver of its SAVNET Agent receives SAV-specific information from other SAVNET routers that act as source entity. Then, its SAVNET Agent processes SAV-specific information provided by other SAVNET routers, its own SAV-specific information, and/or its local routing information to generate SAV rules on corresponding interfaces. As mentioned above, host-facing routers perform SAV filtering on interfaces facing the host network, customer-facing routers perform SAV filtering on interfaces facing the customer network, and AS border routers perform SAV filtering on interfaces facing another AS.

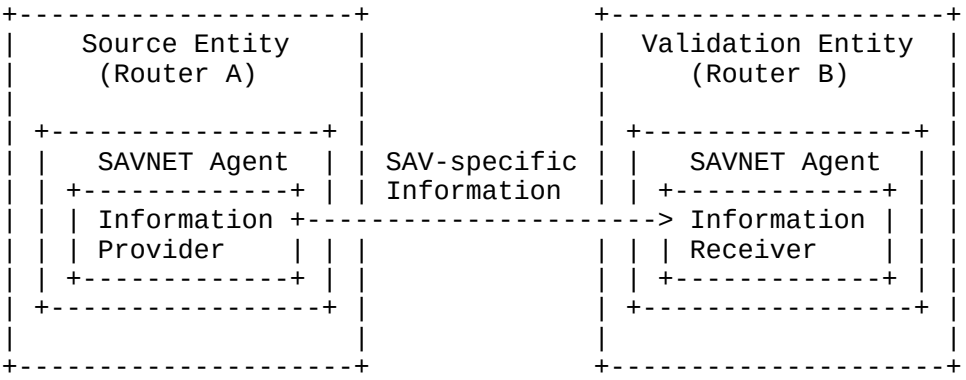

Figure 2: Roles of SAVNET routers

4.3. SAV-specific Information Communication Mechanism

New intra-domain SAV solutions should design a SAV-specific communication mechanism to propagate SAV-specific information from source entity to validation entity. It can be a new protocol or an extension to an existing protocol. This document does not present the details of the protocol design or protocol extensions, but lists necessary features of SAV-specific communication mechanism in the following.

The SAV-specific Information communication mechanism SHOULD define the data structure or format of SAV-specific information, and the operations of communication (such as communication establishment and communication termination). In addition, the mechanism SHOULD require source entity to inform validation entity of the updates of SAV-specific information in a timely manner, so that validation entity can update SAV rules based on the latest information.

In order to ensure the convergence and security of the communication, the session of the SAV-specific communication mechanism SHOULD meet the following requirements:

* The session can be a long-time session or a temporary one, but it SHOULD provide sufficient assurance of transmission reliability and timeliness, so that validation entity can update its SAV rules in time.

* Authentication can be conducted before session establishment. Authentication is optional but the ability of authentication SHOULD be available.

5. SAV-related Information

For intra-domain SAV, both SAV-specific information and local routing information can be used for SAV decisions.

5.1. SAV-specific Information

SAV-specific information is specialized for SAV and thus helps generate more accurate SAV rules. A SAVNET router can obtain its own SAV-specific information based on local routing information, local interface configurations, and/or other local configuration information. In addition, SAVNET routers acting as validation entity can obtain SAV-specific information of other SAVNET routers that act as source entity. By using SAV-specific information provided by other SAVNET routers, the SAVNET router acting as validation entity can generate more accurate SAV rules than solely using its local routing information.

For example, customer-facing routers connected to the same multi-homed customer network can exchange locally-known source prefixes of the customer network through SAV-specific information communication. By processing both SAV-specific information of itself and SAV-specific information of the other customer-facing routers, each of them can identify all prefixes in the customer network and thus avoid improper block in case there is an asymmetric routing. Section 8.1 elaborates on this example.

5.2. Routing Information

Routing information is used for computing packet forwarding rules, which is stored in the router's RIB/FIB. Although it is not specialized for SAV, it is widely used to infer SAV rules in existing uRPF-based SAV mechanisms, such as strict uRPF and loose uRPF [RFC3704]. A SAVNET router acting as validation entity can obtain routing information from its local RIB/FIB to generate SAV rules for some prefixes, when the corresponding SAV-specific information is missing.

6. SAV Rule Generation

Figure 3 shows the SAV rule generation process of the SAVNET router acting as validation entity. The SAV Information Manager of SAVNET Agent consolidates SAV-specific information provided by other routers, SAV-specific information of the router itself, and local routing information into the SAV Information Base. Then, it sends the consolidated information to the SAV Rule Generator. The SAV Rule Generator should preferentially use SAV-specific information to generate SAV rules for specific source prefixes. Local routing information is only recommended when some SAV-specific information is missing.

SAV Information Manager also provides the support of diagnosis. Operators can look up the information in SAV Information Base for monitoring or troubleshooting purpose.

For example, for a host-facing router (or a customer-facing router), it processes SAV-related information to identify prefixes in the host network (or customer network) it connected to, and then generate SAV rules on the interface facing to the host network (or customer network). Data packets coming from that interface will be considered invalid and should be blocked if they use source addresses not belonging to the host network (or customer network). In the incremental/partial deployment scenario when some routers do not deploy SAV-specific information communication mechanism, the host-facing router (or customer-facing router) may not be able to identify all prefixes in the host network (or customer network) through SAV-

specific information. To avoid improper block in this case, the router is recommended to use less strict SAV rules. For example, it can choose to only block packets with non-global or non-routable source addresses by using its local routing information.

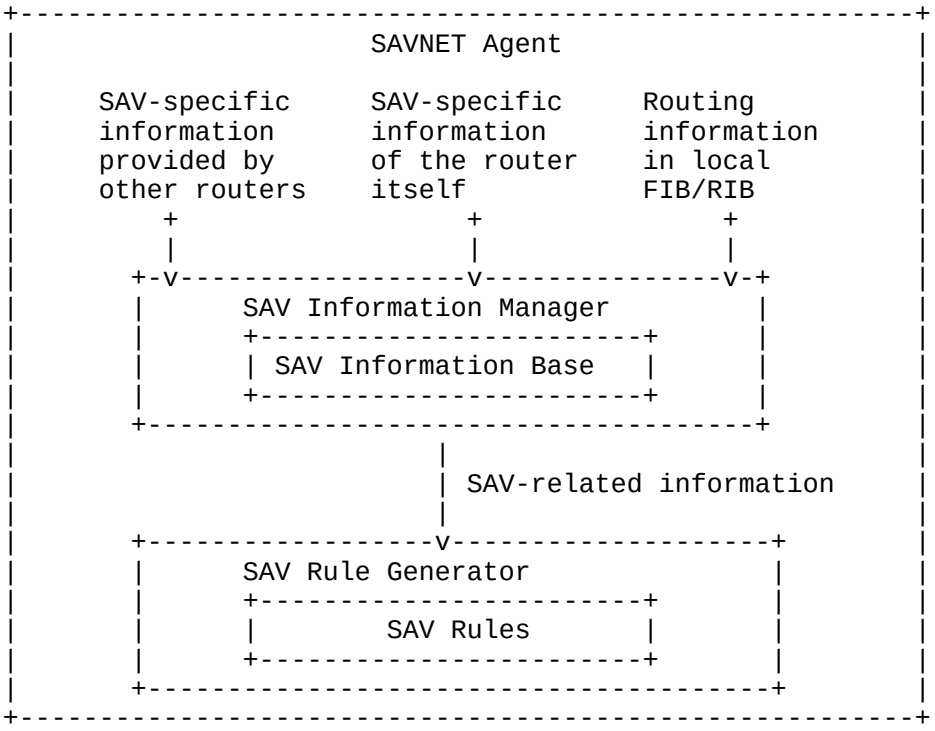

Figure 3: Workflow of SAV rule generation

For an AS border router, it processes SAV-related information to identify prefixes in the local AS, and then generate SAV rules on the interface facing to another AS. Data packets coming from that interface will be considered invalid and should be blocked if they use source addresses belonging to the local AS. In the incremental/ partial deployment scenario, the AS border router may only identify partial prefixes in the local AS through SAV-specific information. In this case, the AS border router can still block data packets with source addresses in learned prefixes.

In addition, if the AS border router also implements inter-domain SAVNET, its intra-domain SAVNET Agent SHOULD send the intra-domain SAV-specific information to its inter-domain SAVNET Agent, helping the inter-domain SAVNET Agent generate inter-domain SAV rules or inter-domain SAV-specific information.

7. Where to deploy intra-domain SAV

A SAVNET router can be a host-facing router, a customer-facing router, an AS border router, or other routers. To reduce deployment overhead and redundant validation, it is not necessary to deploy intra-domain SAV on all intra-domain routers. Future solutions should specify which routers deploy intra-domain SAV and provide incremental benefits when those routers incrementally deploy intra-domain SAV. To this end, this document provides some key recommendations and considerations that should be considered by future solutions.

In general, host-facing routers, customer-facing routers, and AS border routers are vantage points to implement intra-domain SAV. It is not only because these routers are closer to the source and thus will be more effective in identifying and discarding source-spoofed data packets, but also becasue they can clearly determine the directionality of specific source prefixes based on the network topology:

* Host-facing routers (e.g., Router C in Figure 1) generate SAV rules on interfaces facing a host network and only permit incoming data packets that use a source address belonging to the host network.

* Customer-facing routers (e.g., Routers A and B in Figure 1) generate SAV rules on interfaces facing a customer network and only permit incoming data packets that use a source address belonging to the customer network.

* AS border routers (e.g., Routers D or E in Figure 1) generate SAV rules on interfaces facing an external AS and block incoming data packets that use a source address of the local AS.

When only parts of the edge have deployed SAV, every router that has deployed SAV can block spoofing traffic from the connected host network, customer network, or external AS. The local AS is only vulnerable to spoofing traffic entering from parts of the edge where SAV has not been deployed. The network operator can plan the incremental edge deployments by understanding the incremental benefits.

Implementing SAV on other inner routers should be more complicated because many factors will affect the forwarding path from the source to this kind of routers. For example, Traffic Engineering (TE) or Fast Reroute (FRR) is commonly used in an intra-domain network to control the forwarding decisions of routers. To ensure the accuracy of SAV on inner routers, the computation of SAV rules needs to take all factors that will affect forwarding into account.

8. Use Cases

This section uses two use cases to illustrate that intra-domain SAVNET can achieve more accurate and efficient SAV than existing intra-domain SAV mechanisms. The two use cases have already been described in [I-D.ietf-savnet-intra-domain-problem-statement] to show that existing intra-domain SAV mechanisms have problems of improper block or high operational overhead.

8.1. Use Case 1: SAV at Host-facing or Customer-facing Routers

Figure 4 shows an asymmetric routing in a multi-homed host/customer network scenario. Router 1 and Router 2 adopt intra-domain SAV to block spoofing data packets with source addresses not belonging to Network 1 (e.g., a host network or a customer network) receiving from interface '#'.

Network 1 has prefix 10.0.0.0/15 and is connected to two routers (i.e., Router 1 and Router 2) in the intra-domain network. Due to the inbound load balance strategy of Network 1, Router 1 only learns the route to sub prefix 10.1.0.0/16 from Network 1, while Router 2 only learns the route to the other sub prefix 10.0.0.0/16 from Network 1. After that, Router 1 or Router 2 learns the route to the other sub prefix through the intra-domain routing protocol. The FIBs of Router 1 and Router 2 are shown in the figure. Assume Network 1 may send outbound packets with source addresses in sub prefix 10.0.0.0/16 to Router 1 for outbound load balance. The arrows in Figure 4 indicate the direction of traffic.

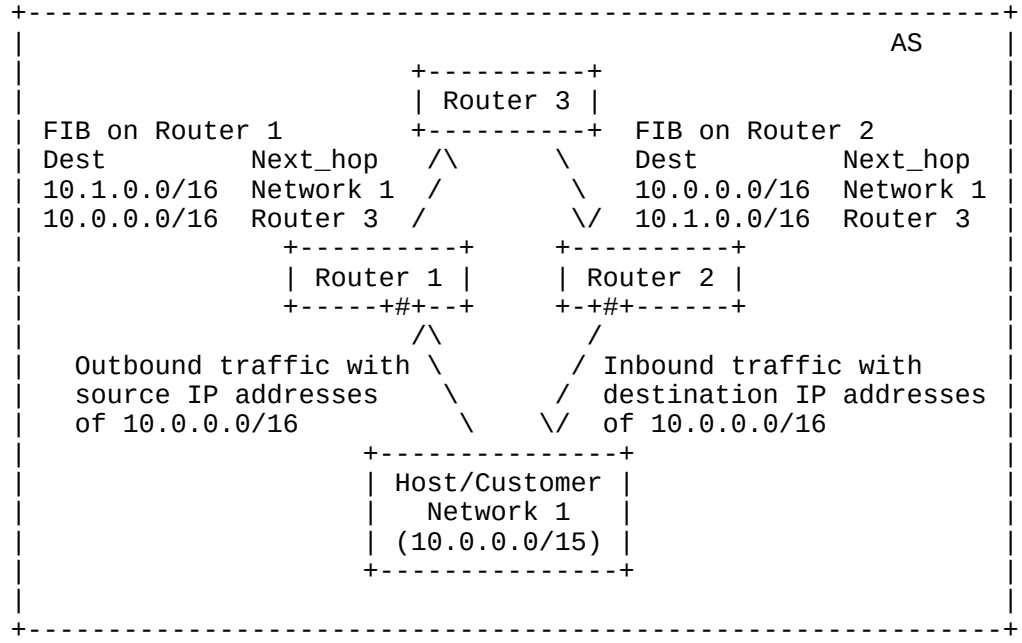

Figure 4: A use case of outbound SAV

In this case, strict uRPF at Router 1 will improperly block legitimate packets with source addresses in prefix 10.0.0.0/16 from Network 1 on interface '#', because it only accepts data packets with source addresses in prefix 10.1.0.0/16 from Router 1's interface '#' according to its local routing information.

If intra-domain SAVNET is implemented in the intra-domain network, Router 2 can inform Router 1 that prefix 10.0.0.0/16 also belongs to Network 1 by providing its SAV-specific information to Router 1. Then, by combining both its own SAV-specific information and SAV-specific information provided by Router 2, Router 1 learns that Network 1 have both prefix 10.1.0.0/16 and prefix 10.0.0.0/16. Therefore, Router 1 will accept data packets with source addresses in prefix 10.1.0.0/16 and prefix 10.0.0.0/16 on interface '#', so improper block can be avoided.

8.2. Use Case 2: SAV at AS Border Routers

Figure 5 shows a scenario of inbound SAV at AS border routers. Router 3 and Router 4 adopt intra-domain SAV to block spoofing data packets with internal source addresses receiving from interface '#'. The arrows in Figure 5 indicate the direction of spoofing traffic.

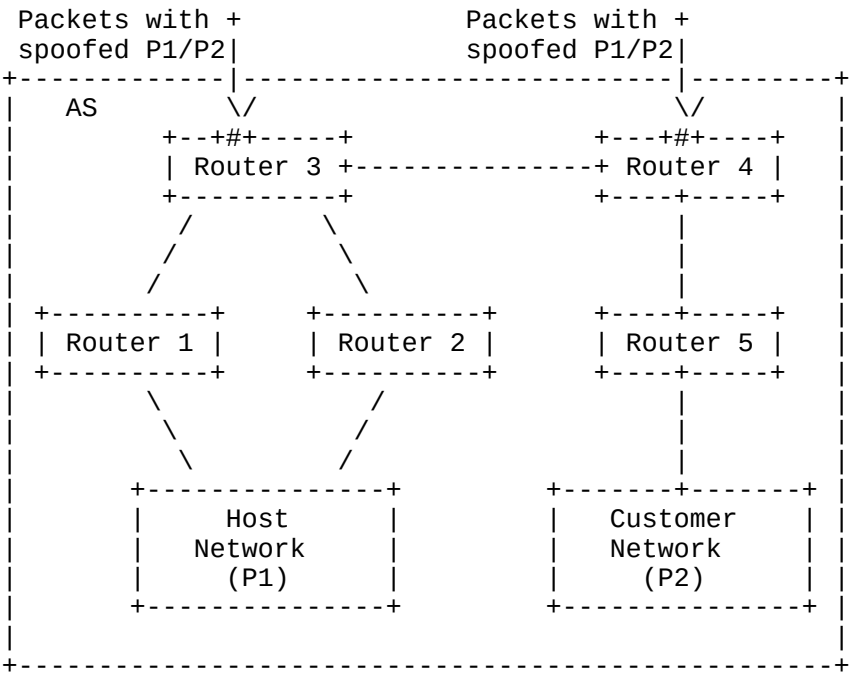

Figure 5: A use case of inbound SAV

If Router 3 and Router 4 deploy ACL-based ingress filtering, the operator needs to manually generate and update ACL rules at Router 3 and Router 4 when internal source prefixes change. The operational overhead of manually maintaining and updating ACL rules will be extremely high, especially when there are multiple inbound validation interfaces '#'.

If intra-domain SAVNET is implemented in the intra-domain network, Router 1, Router 2, and Router 5 will automatically inform Router 3 and Router 4 of prefixes in the host network and customer network by providing SAV-specific information. After receiving SAV-specific information from other routers, Router 3 and Router 4 can identify all internal source prefixes. The SAV-specific information communication will be triggered if topology or prefix related to the host network or customer network changes. For example, if the customer network has a new source prefix P3, Router 5 will inform Router 3 and Router 4 of the new source prefix immediately through SAV-specific information communication mechanism. In this way, Router 3 and Router 4 can automatically generate and update SAV rules on interface '#'.

9. Meeting the Design Requirements of Intra-domain SAVNET

Intra-domain SAVNET architecture is proposed to meet the five design requirements defined in [I-D.ietf-savnet-intra-domain-problem-statement].

9.1. Accurate Validation

In the asymmetric routing scenario shown in Figure 4, the host-facing router (or customer-facing router) cannot identify all prefixes in its host network (or customer network) solely using its local routing information. As a result, existing intra-domain SAV mechanisms (e.g., strict uRPF) solely using local routing information to generate SAV rules will have improper block problems in the case of asymmetric routing.

Intra-domain SAVNET requires routers to exchange SAV-specific information among each other. The SAVNET router can use SAV-specific information provided by other routers as well as its own SAV-specific information to generate more accurate SAV rules. The use case in Figure 4 has shown that intra-domain SAVNET can achieve more accurate SAV filtering compared with strict uRPF in asymmetric routing scenarios.

9.2. Automatic Update

In real intra-domain networks, the topology or prefixes of networks may change dynamically. The SAV mechanism MUST automatically update SAV rules as the network changes. However, ACL-based SAV mechanism requires manual efforts to accommodate to network dynamics, resulting in high operational overhead.

Intra-domain SAVNET allows SAVNET routers to exchange the changes of SAV-specific information among each other automatically. After receiving updated SAV-specific information from source entity, SAVNET routers acting as validation entity can generate and update their SAV rules accordingly. The use case in Section 8.2 has shown that intra-domain SAVNET can achieve automatic update.

9.3. Incremental/Partial Deployment

Although an intra-domain network mostly has one administration, incremental/partial deployment may still exist due to phased deployment or multi-vendor supplement. In phased deployment scenarios, SAV-specific information of non-deploying routers is not available.

As described in Section 6, intra-domain SAVNET can adapt to incremental/partial deployment. To mitigate the impact of phased deployment, it is RECOMMENDED that routers facing the same host/customer network can simultaneously adopt intra-domain SAVNET so that all prefixes in the host/customer network can be identified. For example, in Figure 4, Router 1 and Router 2 are recommended to be upgraded to SAVNET routers together so that the two routers can identify all prefixes in Network 1 and generate accurate SAV rules on interfaces '#'.

In addition, SAVNET routers acting as validation entity are RECOMMENDED to support flexible validation modes and perform SAV filtering gradually to smooth the transition from partial to full deployment:

* SAVNET routers acting as validation entity are RECOMMENDED to support flexible validation modes such as interface-based prefix allowlist, interface-based prefix blocklist, and prefix-based interface allowlist (see [huang-savnet-sav-table]). The first two modes are interface-scale, and the last one is device-scale. Under incremental/partial deployment, SAVNET routers SHOULD take on the proper validation mode according to acquired SAV-specific information. For example, if a customer-facing router can identify all prefixes in its customer network by processing acquired SAV-specific information, an interface-based prefix allowlist containing these prefixes can be used on that customer-facing interface. Otherwise, it should use interface-based prefix blocklist or prefix-based interface allowlist to avoid improper block.

* Validation entity is RECOMMENDED to performed SAV-invalid filtering gradually. The router can first take conservative actions on the validated data packets. That is to say, the router will not discard packets with invalid results in the beginning of deployment. It can conduct sampling action for measurement analysis at first, and then conducts rate-limiting action or redirecting action for packets with invalid results. These conservative actions will not result in serious consequences if some legitimate packets are mistakenly considered invalid, while still providing protection for the network. Finally, filtering action is enabled only after confirming that there are no improper block problems.

9.4. Convergence

When SAV-related information changes, the SAVNET Agent MUST be able to detect the changes in time and update SAV rules with the latest information. Otherwise, outdated SAV rules may cause legitimate data packets to be blocked or spoofing data packets to be accepted.

Intra-domain SAVNET requires routers to update SAV-specific information and update SAV rules in a timely manner. Since SAV-specific information is originated from source entity, it requires that source entity MUST timely send the updated SAV-specific information to validation entity. Therefore, the propagation speed of SAV-specific information is a key factor affecting the convergence. Consider that routing information and SAV-specific information can be originated and advertised through a similar way, SAV-specific information SHOULD at least have a similar propagation speed as routing information.

9.5. Security

Typically, routers in an intra-domain network can trust each other because they would not compromise intra-domain control-plane architectures and protocols.

However, in some unlikely cases, some routers may do harm to other routers within the same domain. Operators SHOULD be aware of potential threats involved in deploying the architecture. Some potential threats and solutions are as follows:

* Entity impersonation.

  - Potential solution: Mutual authentication SHOULD be conducted before session establishment between two entities.

  - Gaps: Impersonation may still exist due to credential theft, implementation flaws, or entity being compromised. Some other security mechanisms can be taken to make such kind of impersonation difficult. Besides, the entities SHOULD be monitored so that misbehaved entities can be detected.

* Message blocking.

  - Potential solution: Acknowledgement mechanisms MUST be provided in the session between a sender and a receiver, so that message losses can be detected.

- Gaps: Message blocking may be a result of DoS/DDoS attack, man-in-the-middle (MITM) attack, or congestion induced by traffic burst. Acknowledgement mechanisms can detect message losses but cannot avoid message losses. MITM attacks cannot be effectively detected by acknowledgement mechanisms.

* Message alteration.

  - Potential solution: An authentication field can be carried by each message so as to ensure message integrity.

  - Gaps: More overhead of control plane and data plane will be induced.

* Message replay.

  - Potential solution: Authentication value can be computed by adding a sequence number or timestamp as input.

  - Gaps: More overhead of control plane and data plane will be induced.

To meet the security requirement, the above security threats SHOULD be considered when designing the new intra-domain SAV mechanism.

10. Data-plane Considerations

This document mainly focuses on SAV rule generation process on control plane, including exchanging SAV-specific information, consolidating SAV-related information, and generating SAV rules. As for data-plane SAV filtering, SAVNET routers check source addresses of incoming data packets against local SAV rules and drop those that are identified as using spoofing source addresses. Therefore, the accuracy of data-plane SAV filtering depends entirely on the accuracy of generated SAV rules. More data-plane considerations can be found in [huang-savnet-sav-table].

11. Manageability Considerations

The architecture provides a general framework for communicating SAV-specific information between routers and generating SAV rules based on SAV-specific information and local routing information. Protocol-independent mechanisms SHOULD be provided for operating and managing SAV-related configurations. For example, a YANG data model for SAV configuration and operation is necessary for the ease of management.

SAV may affect the normal forwarding of data packets. The diagnosis approach and necessary logging information SHOULD be provided. SAV Information Base SHOULD store some information that may not be useful for SAV rule generation but is helpful for management. The SAV-specific information communication mechanism SHOULD have monitoring and troubleshooting functions, which are necessary for efficiently operating the architecture.

12. Privacy Considerations

An intra-domain network is mostly operated by a single organization or company, and the advertised SAV-specific information is used within the network. Therefore, the architecture will not import critical privacy issues in usual cases.

13. IANA Considerations

This document has no IANA requirements.

14. Contributors

Mingqing Huang

Email: huangmq@vip.sina.com

Fang Gao

Email: fredagao520@sina.com

15. Acknowledgements

Many thanks to the valuable comments from: Igor Lubashev, Alvaro Retana, Aijun Wang, Joel Halpern, Jared Mauch, Kotikalapudi Sriram, R端diger Volk, Jeffrey Haas, Xiangqing Chang, Changwang Lin, Xueyan Song, etc.

16. References

16.1. Normative 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-15, 7 April 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-problem-statement-15>.

[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>.

[huang-savnet-sav-table] "General Source Address Validation Capabilities", 2023, <https://datatracker.ietf.org/doc/draft-huang-savnet-sav-table/>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

16.2. Informative References

[RFC5210] Wu, J., Bi, J., Li, X., Ren, G., Xu, K., and M. Williams, "A Source Address Validation Architecture (SAVA) Testbed and Deployment Experience", RFC 5210, DOI 10.17487/RFC5210, June 2008, <https://www.rfc-editor.org/info/rfc5210>.

[RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <https://www.rfc-editor.org/info/rfc7039>.

[RFC7513] Bi, J., Wu, J., Yao, G., and F. Baker, "Source Address Validation Improvement (SAVI) Solution for DHCP", RFC 7513, DOI 10.17487/RFC7513, May 2015, <https://www.rfc-editor.org/info/rfc7513>.

[IPSG] "Configuring DHCP Features and IP Source Guard", January 2016, <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_53_se/configuration/guide/2960scg/swdhcp82.html>.

[cable-verify]
"Cable Source-Verify and IP Address Security", January 2021, <https://www.cisco.com/c/en/us/support/docs/broadband-cable/cable-security/20691-source-verify.html>.

Authors' Addresses

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

Jianping Wu
Tsinghua University
Beijing
China
Email: jianping@cernet.edu.cn

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

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

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