

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 13, 2011

J. Wu
J. Bi
Tsinghua Univ.
M. Bagnulo
UC3M
F. Baker
Cisco
C. Vogt, Ed.
Ericsson
March 12, 2011

Source Address Validation Improvement Framework
draft-ietf-savi-framework-03

Abstract

The Source Address Validation Improvement method was developed to complement ingress filtering with finer-grained, standardized IP source address validation. This document describes and motivates the design of the SAVI method.

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 <http://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 September 13, 2011.

Copyright Notice

Copyright (c) 2011 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	4
2.	Model	4
3.	Deployment Options	6
4.	Scalability Optimizations	8
5.	Reliability Optimizations	10
6.	Mix Scenario	10
7.	Acknowledgment	11
8.	IANA Considerations	11
9.	Security Considerations	11
10.	References	12
10.1.	Informative References	12
10.2.	Normative References	12
	Authors' Addresses	13

1. Introduction

Since IP source addresses are used by hosts and network entities to determine the origin of a packet and as a destination for return data, spoofing of IP source addresses can enable impersonation, concealment, and malicious traffic redirection. Unfortunately, the Internet architecture does not prevent IP source address spoofing [[draft-ietf-savi-threat-scope](#)]. Since the IP source address of a packet generally takes no role in forwarding the packet, it can be selected arbitrarily by the sending host without jeopardizing packet delivery. Extra methods are necessary for IP source address validation, to augment packet forwarding with an explicit check of whether a given packet's IP source address is legitimate.

IP source address validation can happen at different granularity: Ingress filtering [[BCP38](#)], a widely deployed standard for IP source address validation, functions at the coarse granularity of networks. It verifies that the prefix of an IP source address routes to the network from which the packet was received. An advantage of ingress filtering is simplicity: the decision of whether to accept or to reject an IP source address can be made solely based on the information available from routing protocols. However, the simplicity comes at the cost of not being able to validate IP source addresses at a finer granularity, due to the aggregated nature of the information available from routing protocols. Finer-grained IP source address validation would be helpful to enable IP-source-address-based authentication, authorization, and host localization, as well as to efficiently identify misbehaving hosts. Partial solutions [[BA2007](#)] exist for finer-grained IP source address validation, but are proprietary and hence often unsuitable for corporate procurement.

The Source Address Validation Improvement method was developed to complement ingress filtering with standardized IP source address validation at the maximally fine granularity of individual IP addresses: It prevents hosts attached to the same link from spoofing each other's IP addresses. To facilitate deployment in networks of various kinds, the SAVI method was designed to be modular and extensible. This document describes and motivates the design of the SAVI method.

2. Model

To enable network operators to deploy fine-grained IP source address validation without a dependency on supportive functionality on hosts, the SAVI method was designed to be purely network-based. A SAVI instance is located on the path of hosts' packets, enforcing the

hosts' use of legitimate IP source addresses according to the following three-step model:

1. Identify which IP source addresses are legitimate for a host, based on monitoring packets exchanged by the host.
2. Bind a legitimate IP address to a link layer property of the host's network attachment. This property, called a "binding anchor", must be verifiable in every packet that the host sends, and harder to spoof than the host's IP source address itself.
3. Enforce that the IP source addresses in packets match the binding anchors to which they were bound.

This model allows a SAVI instance to be located anywhere on the link to which the hosts attach, hence enabling different locations for a SAVI instance. One way to locate a SAVI instance is in the hosts' default router. IP source addresses are then validated in packets traversing the default router, yet the IP source addresses in packets exchanged locally on the link may bypass validation. Another way to locate a SAVI instance is in a switch between the hosts and their default router. Thus, packets may undergo IP source address validation even if exchanged locally on the link.

The closer a SAVI instance is located to the hosts, the more effective the SAVI method is. This is because each of the three steps of the SAVI model can best be accomplished in a position close to the host:

- o Identifying a host's legitimate IP source addresses is most efficient close to the host, because the likelihood that the host's packets bypass a SAVI instance, and hence cannot be monitored, increases with the distance between the SAVI instance and the host.
- o Selecting a binding anchor for a host's IP source address is easiest close to the host, because many link layer properties are unique for a given host only on a link segment directly attaching to the host.
- o Enforcing a host's use of a legitimate IP source address is most reliable when pursued close to the host, because the likelihood that the host's packets bypass a SAVI instance, and hence do not undergo IP source address validation, increases with the distance between the SAVI instance and the host.

The preferred location of SAVI instances is therefore close to hosts, such as in switches that directly attach to the hosts whose IP source

addresses are being validated.

To make SAVI of better applicability, to handle the scenarios in which hosts are not directly attached is of value. For example, a SAVI instance is attached by a legacy switch which is attached by hosts. However, to take such scenario into consideration requires support in the bind-and-identify model. Considering two scenarios:

- o A legacy switch to which hosts are attaching uses two trunked ports to connect to SAVI switch.
- o STP runs among switches, and a link failure happens.

Both scenarios require more specified design to build correct binding. Thus, for each SAVI solution, the applicability must be specified explicitly.

Besides, in the case of legacy switches, the security level is lower, as there is no full protection for the hosts connected to the legacy server.

3. Deployment Options

The model of the SAVI method, as explained in [Section 2](#), is deployment-specific in two ways:

- o The identification of legitimate IP source addresses is dependent on the IP address assignment method in use on a link, since it is through assignment that a host becomes the legitimate user of an IP source address.
- o Binding anchors are dependent on the technology used to build the link on which they are used, as binding anchors are link layer properties of a host's network attachment.

To facilitate the deployment of the SAVI method in networks of various kinds, the SAVI method is designed to support different IP address assignment methods, and to function with different binding anchors. Naturally, both the IP address assignment methods in use on a link and the available binding anchors have an impact on the functioning and the strength of IP source address validation. The following two sub-sections explain this impact, and describe how the SAVI method accommodates this.

3.1. IP Address Assignment Methods

Since the SAVI method traces IP address assignment packets, it necessarily needs to incorporate logic that is specific to particular IP address assignment methods. However, developing SAVI method variants for each IP address assignment method is alone not sufficient, since multiple IP address assignment methods may co-exist on a given link. The SAVI method hence comes in multiple variants: for links with Stateless Address Autoconfiguration [[rfc4862](#)], for links with DHCP [[rfc2131](#)][[rfc3315](#)], for links with Secure Neighbor Discovery [[rfc3971](#)], for links with IKEv2 [[rfc5996](#)] [[rfc5739](#)] [[rfc5026](#)] and for links that use any combination of IP address assignment methods.

The reason to develop SAVI method variants for each single IP address configuration method, in addition to the variant that handles all IP address assignment methods, is to minimize the complexity of the common case: many link deployments today either are constrained to a single IP address assignment methods or, equivalently from the perspective of the SAVI method, separate IP address assignment methods into different IP address prefixes. The SAVI method for such links can be simpler than the SAVI method for links with multiple IP address assignment methods per IP address prefix.

3.2. Binding Anchors

The SAVI method supports a range of binding anchors:

- o The IEEE extended unique identifier, EUI-48 or EUI-64, of a host's interface.
- o The port on an Ethernet switch to which a host attaches.
- o The security association between a host and the base station on wireless links.
- o The combination of a host interface's link-layer address and a customer relationship in cable modem networks.
- o An ATM virtual channel, a PPPoE session identifier, or an L2TP session identifier in a DSL network.
- o A tunnel that connects to a single host, such as an IP-in-IP tunnel, a GRE tunnel, or an MPLS label-switched path.

The various binding anchors differ significantly in the security they provide. IEEE extended unique identifiers, for example, fail to render a secure binding anchor because they can be spoofed with

little effort. And switch ports alone may be insufficient because they may connect to more than a single host, such as in the case of concatenated switches.

Given this diversity in the security provided, one could define a set of possible binding anchors, and leave it up to the administrator to choose one or more of them. Such a selection of binding anchors would, of course, have to be accompanied by an explanation of the pros and cons of the different binding anchors. In addition, SAVI devices may have a default binding anchor depending on the lower layers. Such a default could be to use switch ports when available, and MAC addresses otherwise. Or to use MAC addresses, and switch ports in addition if available.

4. Scalability Optimizations

The preference to locate a SAVI instance close to hosts implies that multiple SAVI instances must be able to co-exist in order to support large links. Although the model of the SAVI method is independent of the number of SAVI instances per link, co-existence of multiple SAVI instances without further measures can lead to higher-than-necessary memory requirements: since a SAVI instance creates bindings for the IP source addresses of all hosts on a link, bindings are replicated if multiple SAVI instances co-exist on the link. High memory requirements, in turn, increase the cost of a SAVI instance. This is problematic in particular for SAVI instances that are located on a switch, since it may significantly increase the cost of such a switch.

To reduce memory requirements for SAVI instances that are located on a switch, the SAVI method enables the suppression of binding replication on links with multiple SAVI instances. This requires manual disabling of IP source address validation on switch ports that connect to other switches running a SAVI instance. Each SAVI instance is then responsible for validating IP source addresses only on those ports to which hosts attach either directly, or via switches without a SAVI instance. On ports towards other switches running a SAVI instance, IP source addresses are not validated. The switches running SAVI instances thus form a "protection perimeter". The IP source addresses in packets passing the protection perimeter are validated by the ingress SAVI instance, but no further validation takes place as long as the packets remain within, or leave the protection perimeter.

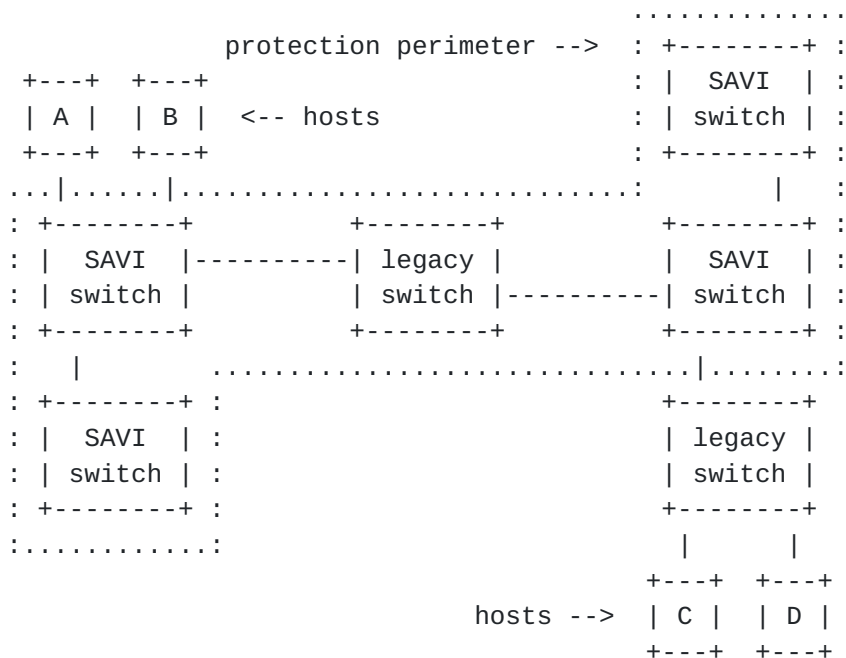


Figure 1: Protection perimeter concept

Figure 1 illustrates the concept of the protection perimeter. The figure shows a link with six switches, of which four, denoted "SAVI switch", run a SAVI instance. The protection perimeter created by the four SAVI instances is shown as a dotted line in the figure. IP source address validation is enabled on all switch ports on the protection perimeter, and it is disabled on all other switch ports. Four hosts, denoted A through D in the figure, attach to the protection perimeter.

In the example of figure Figure 1, the protection perimeter encompasses one of the legacy switches, located in the middle of the depicted link topology. This enables a single, unpartitioned protection perimeter. A single protection perimeter minimizes memory requirements for the SAVI instances because every binding is kept only once, namely, by the SAVI instance that attaches to the host being validated. Excluding the legacy switch from the protection perimeter would result in two smaller protection perimeters to the left and to the right of the depicted link topology. The memory requirements for the SAVI instances would then be higher: since IP source address validation would be activated on the two ports connecting to the legacy switch, the SAVI instances adjacent to the legacy switch would replicate all bindings from the respectively other protection perimeter. The reason why it is possible to include the legacy switch in the protection perimeter is because the depicted link topology guarantees that packets cannot enter the protection

perimeter via this legacy switch. Without this guarantee, the legacy switch would have to be excluded from the protection perimeter in order to ensure that packets entering the protection perimeter undergo IP source address validation.

5. Reliability Optimizations

The explicit storage of legitimate IP addresses in the form of bindings implies that failure to create a binding, or the premature removal of bindings, can lead to loss of legitimate packets. There are three situations in which this can happen:

- o Legitimate IP address configuration packets, which should trigger the creation of a binding in a SAVI instance, are lost before reaching the SAVI instance.
- o A SAVI instance loses a binding, for example, due to a restart.
- o The link topology changes, resulting in hosts to communicate through SAVI instances that do not have a binding for those hosts' IP addresses.

To limit the disruption that missing bindings for legitimate IP addresses can have, the SAVI method includes a mechanism for reactive binding creation based on regular packets. This mechanism supplements the proactive binding creation based on IP address configuration packets. Reactive binding creation occurs when a SAVI instance recognizes excessive drops of regular packets originating from the same IP address. The SAVI instance then verifies whether said IP address is unique on the link. How the verification is carried out depends on the IP address configuration method that the SAVI instance supports: the SAVI method variant for Stateless Address Autoconfiguration and for Secure Neighbor Discovery verifies an IP address through the Duplicate Address Detection procedure. The SAVI method variant for DHCP verifies an IP address through a DHCP Lease Query message exchange with the DHCP server. If verification indicates that the IP address is unique on the link, the SAVI instance creates a binding for the IP address. Otherwise, no binding is created, and packets sent from the IP address continue to be dropped.

6. Mix Scenario

While multiple assignment methods can be used on the same link, the SAVI device may have to deal with a mix of binding discovery methods. If the address prefix used for each assignment method is different,

mix scenario can handle the same as scenario with only one assignment method. If different address assignment methods are used to assign addresses from the same prefix, additional considerations are needed because one binding mechanism may create a binding violating an existing binding from another binding mechanism, e.g., binding from SAVI-FCFS [[savi-fcfs](#)] may violate binding from SAVI-DHCP [[savi-dhcp](#)]. Thus, the collision between different SAVI mechanisms in mix scenario must be handled in case more than one address assignment method is used to assign addresses from the same prefix.

Prioritization relationship between different address assignment methods is used as the basis to solve possible collisions. Current standard documents of address assignment methods have implied the prioritization relationship in general cases. However, considering in some scenarios, default prioritization level may not be quite suitable. Configurable prioritization level should be supported in a document of SAVI solution for the mix scenario.

7. Acknowledgment

The author would like to thank the SAVI working group for a thorough technical discussion on the design and the framework of the SAVI method, as captured in this document, in particular Erik Nordmark, Guang Yao, Eric Levy-Abegnoli, and Alberto Garcia. Thanks also to Torben Melsen for reviewing this document.

This document was generated using the xml2rfc tool.

8. IANA Considerations

This memo asks the IANA for no new parameters.

Note to RFC Editor: This section will have served its purpose if it correctly tells IANA that no new assignments or registries are required, or if those assignments or registries are created during the RFC publication process. From the authors' perspective, it may therefore be removed upon publication as an RFC at the RFC Editor's discretion.

9. Security Considerations

This document only discusses the possible methods to mitigate the usage of forged IP address, but doesn't propose a mechanism that provides strong security for IP address. If binding anchor is not exclusive for each user, or is without strong security, addresses can

still be forged. Besides, the binding may not accord with the address management requirement, which can be more specified for each client. However, given no new protocol is introduced, the improvements are still acceptable. If there is requirement the usage of IP address must be of strong security, the only way is using cryptographic based authentication.

10. References

10.1. Informative References

- [BA2007] Baker, F., "Cisco IP Version 4 Source Guard", IETF Internet draft (work in progress), November 2007.
- [BCP38] Paul, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", [RFC 2827](#), [BCP 38](#), May 2000.

10.2. Normative References

- [[draft-ietf-savi-threat-scope](#)] McPherson, D., Baker, F., and J. Halpern, "SAVI Threat Scope", [draft-ietf-savi-threat-scope-04](#) (work in progress), March 2011.
- [rfc2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [rfc3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [rfc3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), March 2005.
- [rfc4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.
- [rfc5026] Giaretta, G., Kempf, J., and V. Devarapalli, "Mobile IPv6 Bootstrapping in Split Scenario", [RFC 5026](#), October 2007.
- [rfc5739] Eronen, P., Laganier, J., and C. Madson, "IPv6 Configuration in Internet Key Exchange Protocol Version 2 (IKEv2)", [RFC 5739](#), February 2010.
- [rfc5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)",

[RFC 5996](#), September 2010.

[savi-dhcp]

Bi, J., Wu, J., Yao, G., and F. Baker, "SAVI Solution for DHCP", [draft-ietf-savi-dhcp-07](#) (work in progress), November 2010.

[savi-fcfs]

Nordmark, E., Bagnulo, M., and E. Levy-Abegnoli, "FCFS-SAVI: First-Come First-Serve Source-Address Validation for Locally Assigned Addresses", [draft-ietf-savi-fcfs-05](#) (work in progress), October 2010.

Authors' Addresses

Jianping Wu
Tsinghua University
Computer Science, Tsinghua University
Beijing 100084
China

Email: jianping@cernet.edu.cn

Jun Bi
Tsinghua University
Network Research Center, Tsinghua University
Beijing 100084
China

Email: junbi@tsinghua.edu.cn

Marcelo Bagnulo
Universidad Carlos III de Madrid
Avenida de la Universidad 30
Leganes, Madrid 28911
Spain

Email: marcelo@it.uc3m.es

Fred Baker
Cisco Systems
Santa Barbara, CA 93117
United States

Email: fred@cisco.com

Christian Vogt (editor)
Ericsson
200 Holger Way
San Jose, CA 95134
United States

Email: christian.vogt@ericsson.com

