Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 26, 2009

C. Vogt
Ericsson

January 22, 2009

A Solution Space Analysis for First-Hop IP Source Address Validation
draft-ietf-savi-rationale-00.txt

Status of this Memo

Copyright Notice

Abstract

The IETF working group on Source Address Validation Improvements, SAVI, is chartered to design methods for IP source address validation

that complement ingress filtering with finer-grained protection. This document summarizes the discussion in the SAVI working group and design-related conclusions. The purpose of this is two-fold: First, to guide the design process in the working group with written documentation of decisions and their rationale. Second, to provide a measure for assessing the IP source address validation methods that the working group will eventually deliver.

Table of Contents

1. Introduction

While ingress filtering [RFC 2827, BCP 38] provides a way to validate IP source addresses at an aggregated level, there is not yet a standardized mechanism for IP source address validation at a finer granularity. Having a finer granularity would be helpful in a number of situations, including filtering traffic from customer interfaces implemented as ports in an IP-aware switch or a router, or general improvements in filtering accuracy in enterprise networks. Depending on the situation, there may be a requirement for blocking spoofed packets or merely logging packets that appear to be spoofed.

Partial solutions exist to prevent hosts from spoofing the IP source address of another host in the same IP link (e.g., the "IP source guard"), but are proprietary. The purpose of the IETF working group on Source Address Validation Improvements, SAVI, is to standardize methods that prevent hosts attached to the same IP link from spoofing each other's IP addresses. These methods are to complement ingress filtering with finer-grained protection [TAXO].

This document summarizes the discussion in the SAVI working group and design-related conclusions. The purpose of this is two-fold: First, to guide the design process in the working group with written documentation of decisions and their rationale. Second, to provide a measure for assessing the IP source address validation methods that the working group will eventually deliver.

2. Lower-Layer Binding Anchor

Since the SAVI charter prohibits host changes, a SAVI device will necessarily have to bind IP addresses to a property of layers below IP. This is because, in the absence of host changes, properties of lower layers are the only reasonably trustworthy information about a packet sender that shows up in all packets. The question hence is which lower-layer properties, or lower-layer "binding anchors", are most appropriate for this purpose. Depending on the lower layers, the available options are the following:

- The IEEE extended unique identifier, EUI-48 or EUI-64, of a host's interface.
- The port on an Ethernet switch to which a host attaches.
- The security association between a host and the base station on wireless links.

- The combination of a host interface's link-layer address and a customer relationship in cable modem networks.

- An ATM virtual channel, a PPPoE session identifier, or an L2TP session identifier in a DSL network.

- 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 options, of course, differ significantly in the scurity they provide. IEEE extended unique identifiers, for example, fail to render a secure binding anchor because they can be spoofed without much 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.

One possible approach would be to 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. EIn 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. EOr to use MAC addresses, and switch ports in addition if available.

3. Packet Classification

The prerequisite that a SAVI solution should be complementary to ingress filtering, and not substitute it, implies that SAVI should not validate packets that are forwarded by routers. This calls for a method for SAVI to classify first-hop packets from forwarded packets (where "first-hop packets" are transmitted by the originating host, and "forwarded packets" are relayed by a router). Techniques to achieve such packet classification can be divided into the following classes:

1. Packets are classified based on whether or not their source address is from an on-link subnet prefix.

2. Packets are classified based on whether or not the sending node is an authorized router.

Both classes of packet classification techniques have pros and cons. An advantage of class (1) is that the configuration of SAVI devices with the necessary packet classification information is in many cases simpler: SAVI devices that are colocated with a router have direct

access to on-link subnet prefixes because routers need to be aware of the on-link subnet prefixes themselves. Furthermore, SAVI devices can learn on-link subnet prefixes by listening to DHCP messages and, in IPv6, to Router Advertisement messages. This enables auto-configuration of SAVI devices that are implemented on a switch. With class (2), similar auto-configuration is possible only with SeND because a router can then be securely identified based on its verifiable Router Advertisement messages. There is no way for a SAVI device to automatically and securely identify a router if plain Neighbor Discovery is used. Class (2) therefore requires pre-configuration of SAVI devices with information about local routers if plain Neighbor Discovery is used.

Of course, the auto-configuration of SAVI devices via DHCP messages or Router Advertisement messages requires that the complete set of on-link subnet prefixes is announced in these messages. It is insufficient where DHCP is not used and no Router Advertisement messages are sent, or where not all on-link subnet prefixes are reveiled in DHCP messages or Router Advertisement messages. SAVI devices then require additional sources for on-link subnet prefix information. For example, on-link subnet prefixes that are manually configured into hosts or routers have to be configured also into SAVI devices.

On the other hand, a disadvantage of class (1) is that SAVI may erroneously discard legitimate packets. This happens when a host sends a packet to a neighbor via the local router instead of sending it to the neighbor directly. The packet is then dropped when forwarded by the local router because it is considered a first-hop packet based on the subnet prefix of its source address. With class (2), the SAVI device would not validate, and hence not drop, the packet given that it is coming from the router. This issue with class (1) can be mitigated if local routers use Redirect messages to enforce direct neighbor-to-neighbor communications.

One conclusion from the above could be that class (2) should only be used when SeND is available. And in such a case, class (1) could be omitted. This would mean that, with plain Neighbor Discovery, class (1) would be used exclusively.

4. Acknowledgment

This document is a resume of the discussions of the IETF working group on Source Address Validation Improvements. The author would therefore like to thank the working group members for their valuable contributions, especially Marcelo Bagnulo, Fred Baker, Jun Bi, Wojciech Dec, Paul Ferguson, David Miles, Erik Nordmark, Pekka

Savola, Dave Thaler, Guang Yao, Mark Williams, Jianping Wu, Dong Zhang, and Lixia Zhang, in alphabetical order.

This document was generated using the xml2rfc tool.

5. Normative References

[TAXO] Vogt, C. and J. Arkko, "SAVI Design Taxonomy and Analysis", July 2008.

Author's Address

Christian Vogt
Ericsson Research, NomadicLab
Office 551
Hirsalantie 11
02420 Jorvas
Finland

Email: christian.vogt@ericsson.com