
Workgroup: SAVNET Working Group
Internet-Draft: draft-cui-savnet-anti-ddos-01
Published: 13 March 2023
Intended Status: Informational
Expires: 14 September 2023
Authors: Y. Cui J. Wu L. Hui
Tsinghua University Tsinghua University Zhongguancun Laboratory
L. Zhang Y. Hu
Zhongguancun Laboratory Zhongguancun Laboratory

SAV-based Anti-DDoS Architecture

Abstract

Existing SAV schemes can not effectively defend against IP Spoofing DDoS under incremental deployment. This document proposes SAV-D, an SAV-honeynet based distributed defense architecture to enhance SAV's defense. The main idea of SAV-D is to collect and aggregate more threat data from existing SAV devices and then distribute crucial knowledge to widespread devices, thus significantly expanding defense across the entire network.

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 14 September 2023.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions

with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- 1. Introduction
 - 1.1. Requirements Language
- 2. Problem Statement
- 3. SAV-D Architecture
 - 3.1. SAV Controller
 - 3.2. SAV Device
 - 3.3. Legacy Device
 - 3.4. Victims' Defense
 - 3.5. Connection Example
- 4. Workflow
- 5. Scalability
- 6. IANA Considerations
- 7. Security Considerations
- 8. References
 - 8.1. Normative References
 - 8.2. Informative References
- Acknowledgements
- Authors' Addresses

1. Introduction

Distributed Denial-of-Service (DDoS) attacks have been a persistent cyber threat, where IP spoofing DDoS is one of the major contributors. Amplification DDoS typically exploit IP spoofing to generate large volumes of traffic with small requests, allowing attackers to overwhelm the target's resources while evading detection. Some other DDoS attacks (e.g., TCP SYN Flooding [RFC4987]) also forge source IP addresses in order to drain the target's resources.

To eliminate IP spoofing, several Source Address Validation (SAV) schemes have been proposed, such as SAVI[RFC7039], uRPF[RFC3704] and EFP-uRPF[RFC8704]. However, the defense effectiveness of current SAV schemes highly depends on the SAV devices' deployment ratio. A large number of spoofed packets can only be prevented with a significantly high deployment ratio, but the incremental deployment process is often slow. According to CAIDA's Spoofer Project[CAIDA], 28.7% of IPv4 autonomous systems (excluding NAT), and 34.3% of IPv6 autonomous systems are still spoofable by March 2023. This indicates a limited SAV deployment, thus the defense effectiveness.

In the above context, this document offers an SAV-based anti-DDoS architecture (SAV-D) that incorporates the following advances.

- SAV-honeynet based threat data collection. Each SAV device functions as a honeypot that does not directly drop spoofed packets but instead records the spoofing characteristics and sends them to a centralized control plane.
- Collaborative defense with both SAV and non-SAV devices. The control plane detects ongoing attacks and generates filtering rules. These rules are then distributed to both SAV and non-SAV devices along the attack paths to manipulate malicious traffic.
- Threat information sharing with the victim-end. The control plane shares attack detection information and IP blocklists with victim-end defense systems to assist their mitigations.

Through the mechanisms of honeynet, data aggregation and distribution, SV-D can fully leverage the value of SAV devices and threat data, resulting in a significant defense improvement.

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.

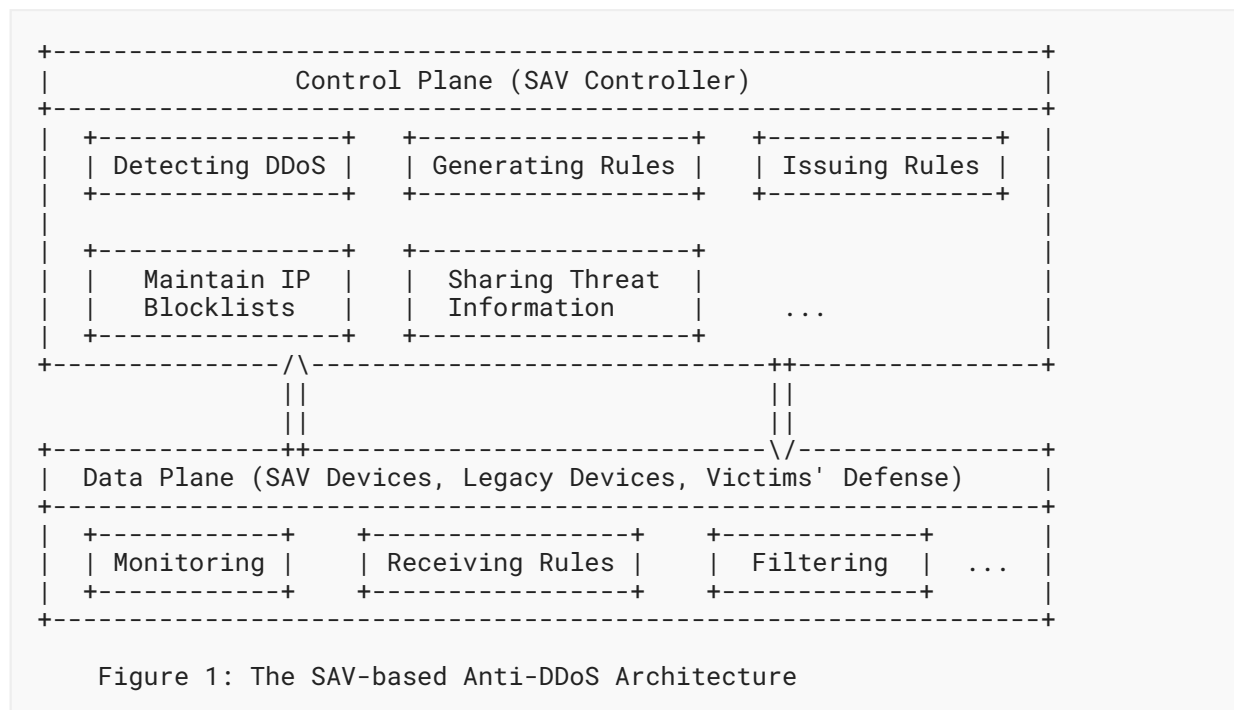
2. Problem Statement

The effectiveness of existing SAV schemes highly relies on the deployment ratio of devices, which is currently limited. Adversaries often actively test their bots for plausibility, packet loss, and amplification benefits. This testing can force the bots to migrate from SAV domains to non-SAV domains, resulting in fewer spoofed packets being blocked by SAV devices. Additionally, uRPF and EFP-uRPF have issues with filtering accuracy in certain scenarios. Some managers may hesitate to enable SAV due to the probability of filtering errors. Moreover, SAV can prevent spoofed packets from being sent out, but it cannot provide protection for the deployers. The lack of direct benefits may also impede the deployment process. In this context, there is a strong need to improve the defense capabilities of current SAV practices.

To achieve the goal, it is essential to consider the following limitations. Firstly, due to the attack testing, directly dropping spoofed packets can reduce the possibility of capturing threat data. Secondly, in amplification DDoS, the reflected packets sent to victims have the authentic src-IP,

making them unfilterable by SAV devices. Lastly, although today's SAV mechanism can filter spoofed packets at local devices, the important threat information they provide has yet to be fully utilized. If victims were made aware of the type of spoofing traffic targeting them, they could execute faster and more accurate countermeasures.

3. SAV-D Architecture



The proposed SAV-D is shown in Figure 1. It introduces a centralized control plane (i.e., the controller) that connects SAV devices, legacy devices, and victims' defense systems. The controllers can collect spoofing characteristics from widespread SAV devices (as honeypots) and aggregate them for further analysis. From a whole viewpoint, the controller can detect ongoing attacks and generate filtering rules for both SAV and non-SAV devices. These rules will be distributed to corresponding devices to perform filtering. Additionally, the controller will share the attack information with the victims' defense system to assist in their defense operations.

3.1. SAV Controller

The controller is a logical entity that can be implemented as a distributed or centralized cluster system. The placement of controllers may take several factors into consideration, including latency, resiliency, and load balancing to connected devices.

- To collect spoofing information, the controller will passively receive the data sent from the certified SAV devices. The collected spoofing information should include but not limited to timestamp, 5-tuple (i.e., src-IP, dst-IP, src-port, dst-port, and protocol), TCP flag, packet size, and amounts. This information will be readily stored in a database for further analysis.

- To analyze the aggregated statistics, the controller retrieves the spoofing information periodically (e.g., every 10 seconds). The spoofed packets are analyzed based on their src-IP to detect reflection attacks with certain algorithms. A large volume of spoofed packets using a specific protocol (e.g., NTP, DNS) is a clear indication that the src-IP is being targeted by reflection attacks. The detection results include the attack target, type, duration, malicious IP lists, etc.
- Generating filtering rules based on detection results is a straightforward process. Before the reflection, the filtering rules are based on src-IP and ports. After reflection, the src-IP is the server's address, and the dst-IP is the victim's address. Considering the reflected packets are often much larger than legitimate packets, filtering rules could be generated based on dst-IP, ports, and packet size.
- Communicating with relevant devices consists of two folds. One fold is distributing filtering rules to SAV and legacy devices and receiving feedback from SAV devices. The other fold is to provide the victim's defense system with attack detection information, which is essential to efficiently stop the attack traffic.

3.2. SAV Device

The SAV devices refer to routers or switches that are capable of validating the source IP address, including SAVI, uRPF, etc. Compared to simply dropping spoofed packets, SAV devices are required to selectively allow spoofed packets through if they do not match the filtering rules. This mechanism can be considered as a SAV-honeynet that records threat data related to spoofing.

- The SAV device must register it to the controller when being installed, in which a unique identification number and other information (e.g., location, management IP address) may be needed. Whenever a spoofed packet is detected, the SAV device will record its timestamp, 5-tuple, TCP flag, packet size, and so on. However, only if the spoofed packet matches existing filtering rules, will the packet be dropped. After a certain interval, the recorded data will be compressed and sent to the controller.
- Modern devices are generally capable of filtering based on packet length and counting the number of filtered packets. Upon receiving filtering rules from the controller, the SAV device must install them into its data plane. The SAV device also needs to record the number of packets filtered by each rule. If a rule filters no packet during some periods, the rule will be automatically removed to save the rule's space.

3.3. Legacy Device

The commercial routers that are widely deployed in production are considered to be legacy devices. Access Control List (ACL) is universally supported in today's routers for packet filtering. Legacy devices can achieve extensive filtering by simply connecting their management interface to the controller and receiving the rules. Since ACLs may vary across legacy devices, filtering rules must be adapted to meet the specific requirements of each device. The legacy routers can join the SAV-D system by registering it to the controller with information similar to the SAV router. Once registered, the legacy routers can receive the filtering rules from the controller in a safe and trusted channel. These rules will be installed into the data plane. Similar to SAV devices, if a rule filters no packet during some period, the rule will be automatically removed.

3.4. Victims' Defense

The SAV deployers can request access to the attack detection information related to themselves. The information includes various details such as the attack target, type, duration, and malicious IP lists. These details can serve as auxiliary signals to boost the detection time. In addition, SAV-D can provide real-time updated IP blocklists, which can be efficiently used for blocking malicious traffic. In an ideal situation, the defense system could provide an interface to directly receive the information and automatically perform corresponding filtering policies. This mechanism could improve the effectiveness of DDoS defense and incentivize more SAV deployment.

3.5. Connection Example

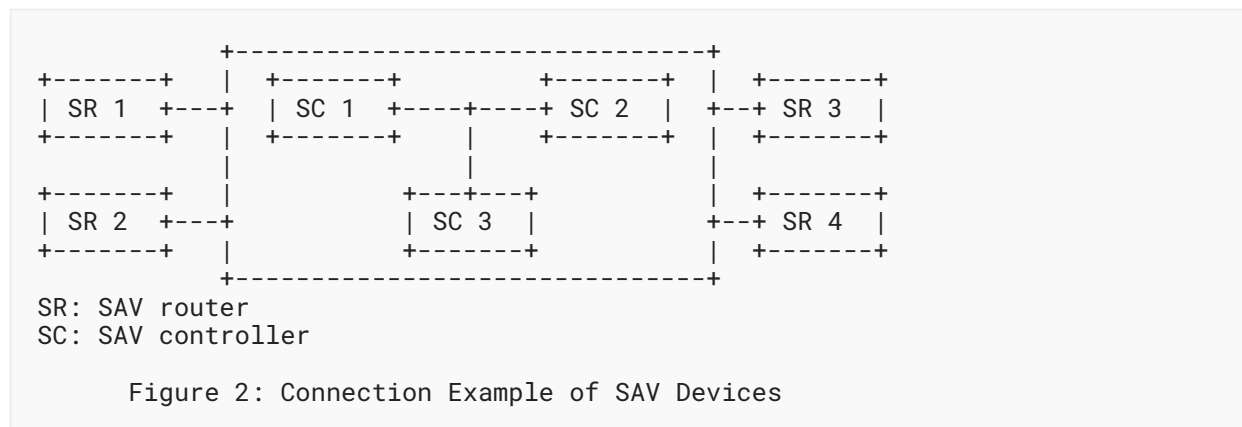


Figure 2 depicts a connection example of SAV-D system. There are SAV routers distributed throughout the network, and they **MUST** communicate with the SAV controller in order to collaborate. This document suggests that each SAV router stores several records of the SAV controller for backup. Each SAV router **MUST** try to connect to its nearest SAV controller at all times. If the SAV router loses contact with the present controller, it **MUST** seek the next closest controller. Such a mechanism can assist SAV routers in maintaining connections to the best of their abilities.

The SAV controller appears as a single entity to the external. Realizing the full functionality of the SAV controller **MAY** require many computing and storage resources. As a result, the SAV controller can be built as clustered or distributed servers, where consistency and scalability are the primary concerns. Each SAV controller can communicate with many SAV routers and perform the corresponding functions.

4. Workflow

The proposed SAV-D architecture can collaboratively defend the IP spoofing DDoS in a distributed pattern. The typical procedures are described as follows.

- (i). The SAV routers validate and record the characteristics of spoofed packets, and periodically send this data to the logically centralized controller, where the global spoofing information is aggregated.
- (ii). Based on the aggregated statistics, the controller can accurately detect whether there are ongoing IP spoofing attacks with the help of predefined algorithms.
- (iii). Based on the detection results, the controller can generate defense policies for both SAV and non-SAV devices. The policies mainly involve filtering rules on 5-tuple and packet size.
- (iv). For detected attacks, the defense policies will be distributed to all SAV and legacy devices. Moreover, the detection results will also be sent to the victim's defense system.
- (v). The filtering rules will be installed on relevant devices to block the malicious packets. If a rule filters no packet during some period, the rule will be automatically removed.

5. Scalability

When there are large amounts of devices introduced into the SAV-D, the control plane could be implemented with hierarchical structure, where multiple sub-level controllers are in charge of the devices inside AS domains. The single top-level controller can exchange information (i.e., IP spoofing statistics and filtering rules) with these sub-level controllers. Additionally, a large number of attacks and filtering rules could introduce another scalability problem. One possible solution is to prioritize the mitigations of these attacks, where severe attacks will be tackled first so that the number of filtering rules will be limited to moderate scope.

6. IANA Considerations

This document includes no request to IANA.

7. Security Considerations

Adversaries may send forged IP spoofing statistics to the control plane or send forged filtering rules to SAV and legacy devices, which could cause severe harm to legitimate traffic. To avoid this situation, the information transmissions of SAV-D could be encrypted with certification. There could also be attacks directly on the SAV-D controllers. As common systems, security systems (e.g., firewalls) are essential to protect the controllers. In addition, hot-standby controllers can also significantly improve security and availability.

8. References

8.1. Normative References

[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/rfc/rfc3704>>.

[RFC8704] Sriram, K., Montgomery, D., and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding", BCP 84, RFC 8704, DOI 10.17487/RFC8704, February 2020, <<https://www.rfc-editor.org/rfc/rfc8704>>.

[RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", RFC 4987, DOI 10.17487/RFC4987, August 2007, <<https://www.rfc-editor.org/rfc/rfc4987>>.

[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/rfc/rfc7039>>.

[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/rfc/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/rfc/rfc8174>>.

8.2. Informative References

[CAIDA] "State of IP Spoofing", March 2023, <<https://spoofer.caida.org/summary.php>>.

Acknowledgements

TBD

Authors' Addresses

Yong Cui

Tsinghua University

Beijing, 100084

China

Email: cuiyong@tsinghua.edu.cn

URI: <http://www.cuiyong.net/>

Jianping Wu

Tsinghua University

Beijing, 100084

China

Email: jianping@cernet.edu.cn

Linbo Hui

Zhongguancun Laboratory
Beijing, 100094
China
Email: huilb@zgclab.edu.cn

Lei Zhang

Zhongguancun Laboratory
Beijing, 100094
China
Email: zhanglei@zgclab.edu.cn

Yannan Hu

Zhongguancun Laboratory
Beijing, 100094
China
Email: huyn@zgclab.edu.cn