

S

Source Address Validation



Maciej Korczyński and Yevheniya Nosyk
University Grenoble Alpes, CNRS, Grenoble
INP, LIG, Grenoble, France

Definitions

Source address validation (SAV) is a standard formalized in RFC 2827 aimed at discarding packets with spoofed source IP addresses. The absence of SAV has been known as a root cause of reflection distributed denial-of-service (DDoS) attacks.

Outbound SAV (oSAV): filtering applied at the network edge to traffic coming from inside the customer network to the outside.

Inbound SAV (iSAV): filtering applied at the network edge to traffic coming from the outside to the customer network.

Background

The Internet relies on IP packets to enable communication between hosts with the destination and source addresses specified in packet headers. However, there is no packet-level authentication mechanism to ensure that the source address has not been altered (Beverly et al. 2009). The modification of a source IP address is referred to as “IP spoofing.” It results in the anonymity

of the sender and prevents a packet from being traced to its origin. This vulnerability has been leveraged to launch distributed denial-of-service (DDoS) attacks that can be made even more effective using reflection (Beverly and Bauer 2005). Because it is not possible in general to prevent packet header modification, concerted efforts have been undertaken to prevent spoofed packets from reaching potential victims. This goal can be achieved by filtering packets at the network edge, formalized in RFC 2827, and called *source address validation* (SAV) (Senie and Ferguson 2000).

The RFC defined the notion of ingress filtering—discarding any packets with source addresses not following filtering rules. This operation is the most effective when applied at the network edge (Senie and Ferguson 2000). RFC 3704 proposed different ways to implement SAV including static access control lists (ACLs) and reverse path forwarding (Baker and Savola 2004).

Packet filtering can be applied in two directions: *inbound* to the customer’s network from outside (Korczyński et al. 2020a) and *outbound* from the customer to outside (Senie and Ferguson 2000). The lack of SAV in any of these directions may result in different security threats.

Attackers benefit from the absence of oSAV to launch DDoS attacks, in particular, reflection attacks. Adversaries make use of public services prone to amplification (Rossow 2014), such as open DNS resolvers or NTP servers, to which they send requests on behalf of their victims by

spoofing their source IP addresses. The victim is then overloaded with the traffic coming from the services rather than from the botnet controlled by the attacker. In this scenario, the origin of the attack is not traceable. One of the most successful attacks against GitHub resulted in traffic of 1.35 Tbps: attackers redirected Memcached responses by spoofing their source addresses (Kottler 2018). In such scenarios, spoofed source addresses of the victims are usually globally routable IPs. In some cases, to impersonate an internal host, a spoofed IP address may be from the inside target network, which reveals the absence of iSAV (Baker and Savola 2004).

Pretending to be an internal host reveals information about the inner network structure, such as the presence of closed DNS resolvers that resolve only on behalf of clients within the same network. The absence of iSAV may have serious consequences when combined with the NXDOMAIN attack, also known as the water torture attack (Luo et al. 2018), or the recently discovered NXNSAttack (Shafir et al. 2020). Both attacks enable denial of service against both recursive resolvers and authoritative servers.

The possibility of impersonating another host on the victim network can also assist in the zone poisoning attack (Korczyński et al. 2016). A DNS server, authoritative for a given domain, may be configured to accept so-called nonsecure DNS dynamic updates from hosts (e.g., a DHCP server) on the same network (Vixie et al. 1997). Therefore, sending a single spoofed UDP packet from the outside with an IP address of that host will modify the content of the zone file (Korczyński et al. 2016). The attack vector can be used to hijack the domain name. Another way to target closed resolvers is to perform DNS cache poisoning (Kaminsky 2008). An attacker can send a spoofed DNS A request for a specific domain to a closed resolver, followed by forged replies before the arrival of the response from the genuine authoritative server. In this case, the users who query the same domain will be redirected to where the attacker specified until the forged DNS entry reaches its time to live (TTL).

Despite the knowledge of the abovementioned attack scenarios and the costs of the damage

they may incur, it was shown that SAV is not yet widely deployed. Lichtblau et al. surveyed 84 network operators to learn whether they deployed SAV and what challenges they faced (Lichtblau et al. 2017). The reasons for not performing packet filtering included incidentally filtering out legitimate traffic, equipment limitations, and lack of a direct economic benefit. In the case of outbound SAV, the compliant network cannot become an attack source but can be attacked itself. Therefore, oSAV suffers from misaligned economic incentives: a network operator that adopts oSAV incurs the cost of deployment, while the security profits benefit all other networks (Lone et al. 2020). On the other hand, performing inbound SAV protects networks from direct threats, which is beneficial from an economic perspective.

Application

Given the prevalent role of IP spoofing in cyberattacks, there is a need to estimate the level of SAV deployment by network providers. Increasing the visibility of the networks that allow spoofing leads to a decrease in the information asymmetry between network operators, their peers, and customers and thus may strengthen the economic incentives for the adoption of SAV.

Table 1 summarizes the methods proposed to infer SAV deployment. They differ in terms of the filtering direction (iSAV versus oSAV) whether they infer the presence or absence of SAV, whether measurements can be done remotely or on a vantage point inside the tested network is required, and if the method relies on existing network misconfigurations.

The Closed Resolver project (Korczyński et al. 2020a,b) aims at mitigating the problem of inbound IP spoofing. They identify closed and open DNS resolvers that accept spoofed requests coming from the outside of their network. The proposed method is remote and does not rely on existing misconfigurations. It provides the most complete picture of iSAV deployment by network providers and covers over 55% IPv4 and 27% IPv6 ASes (Korczyński et al.

Source Address Validation, Table 1 Methods to infer deployment of SAV

Method	SAV direction	Presence/absence	Remote	Relies on misconfigurations
Closed resolver (Korczyński et al. 2020a)	iSAV	Both	Yes	No
Spoofers (Beverly and Bauer 2005)	oSAV/iSAV	Both	No	No
Forwarder-based (Kührer et al. 2014)	oSAV	Absence	Yes	Yes
Traceroute loops (Lone et al. 2017)	oSAV	Absence	Yes	Yes
Spoofers-IX (Müller et al. 2019)	oSAV	Both	No	No

2020c). It reveals that the great majority of ASes are fully or partially vulnerable to inbound spoofing. Recently, Deccio et al. (2020) used a similar approach to the methodology proposed by Korczyński et al. (2020a,b,c).

The Spoofers project (Beverly and Bauer 2005; Beverly et al. 2009; Luckie et al. 2019) deploys a client-server infrastructure mainly based on volunteers and “crowdworkers” hired for one study through five crowdsourcing platforms (Lone et al. 2018) that run the client software from inside a network. The active probing client sends both unspoofed and spoofed packets to the Spoofers server either periodically or when it detects a new network. The server inspects received packets (if any) and analyzes whether spoofing is allowed and to what extent (Beverly et al. 2009). This approach identifies the absence and the presence of SAV in both directions. The results obtained by the Spoofers project provide the most confident picture of the deployment of oSAV and have covered tests from 7,915 ASes since 2015 (Spoofers Project 2020). However, those that are not aware of this issue or do not deploy oSAV are less likely to run Spoofers on their networks.

A more practical approach is to perform such measurements remotely. Kührer et al. (2014) scanned for open DNS resolvers, as proposed by Mauch (2013), to detect the absence of outbound SAV. The method leverages the misconfiguration of forwarding resolvers and is referred to as *forwarder-based*. The misbehaving resolver forwards a request to a recursive resolver with either not changing the packet source address to its own address or by sending back the response to the client with the source IP of the recursive resolver. Misconfigured forwarders revealed 2,692 ASes that are fully or partially vulnerable to outbound spoofing.

Lone et al. (2017) proposed another method that does not require a vantage point inside a tested network. When packets are sent to a customer network with an address that is routable but not allocated, this packet is sent back to the provider router without changing its source IP address. The packet, having the source IP address of the machine that sent it, should be dropped by the router because the source IP does not belong to the customer network. The method detected 703 ASes not deploying oSAV.

Finally, while the abovementioned methods rely on actively generated (whether spoofed or not) packets, Müller et al. (2019) passively observed and analyzed inter-domain traffic exchanged between networks at a large IXP taking into account AS business relationships, asymmetric routing, and traffic engineering.

Open Problems and Future Directions

Although the Internet community has developed technical solutions to mitigate the spoofing vulnerability and a variety of methods to estimate the level of SAV deployment by network providers, its deployment remains low. Lack of a direct economic benefit in case of deploying oSAV remains one of the primary problems preventing providers from applying the existing technical standards (Lichtblau et al. 2017). This failure is referred to as *negative externality*: network operators do not invest in implementing the security standard while imposing economic costs on other networks that are victims of attacks using IP spoofing (Luckie et al. 2019).

The deployment of iSAV does not suffer from misaligned economic incentives and protects the provider network that deploys the

standard rather than other networks. Interestingly, SAV for outbound traffic turned out to be more deployed than inbound at the AS level among network operators committed to the Mutually Agreed Norms for Routing Security regulations (MANRS 2020) initiative (Luckie et al. 2019; Korczyński et al. 2020c). At the time of writing, 515 ASes are its signatories. MANRS recommends its members to implement SAV in their networks “to prevent packets with an incorrect source IP address from entering or leaving the network” (Korczyński et al. 2020c). One possible explanation for the higher deployment of oSAV among MANRS members is that the absence of outbound packet filtering gained widespread attention since it is the reason for reflection DDoS attacks. Under these circumstances, the SAV of inbound traffic remained neglected or overlooked by network operators.

“Naming and shaming” of network operators appeared to be a weak form of incentive (Luckie et al. 2019) for deploying oSAV. Luckie et al. (2019) consider several potential future scenarios, including liability associated with attacks originating from their networks or different types of regulations, including governmental initiatives. Finally, long-term efforts taken by the research community to measure and notify noncompliant operators such as the Spoofer project for oSAV and the Closed Resolver project for iSAV may significantly contribute to improving the overall deployment of the standard.

References

- Baker F, Savola P (2004) Ingress filtering for multihomed networks. RFC 3704. <https://rfc-editor.org/rfc/rfc3704.txt>
- Beverly R, Bauer S (2005) The spoofer project: inferring the extent of source address filtering on the internet. In: USENIX Steps to Reducing Unwanted Traffic on the Internet Workshop
- Beverly R, Berger A, Hyun Y, claffy k (2009) Understanding the efficacy of deployed internet source address validation filtering. In: Internet Measurement Conference, ACM
- Deccio C, Hilton A, Briggs M, Avery T, Richardson R (2020) Behind closed doors: A network tale of spoofing, intrusion, and false DNS security. In: Internet Measurement Conference, ACM
- Kaminsky D (2008) It’s the end of the cache as we know it. <https://www.slideshare.net/dakami/dmk-bo2-k8>
- Korczyński M, Król M, van Eeten M (2016) Zone poisoning: the how and where of non-secure DNS dynamic updates. In: Internet Measurement Conference, ACM
- Korczyński M, Nosyk Y, Lone Q, Skwarek M, Jonglez B, Duda A (2020a) Don’t forget to lock the front door! Inferring the deployment of source address validation of inbound traffic. In: Passive and active measurement. Springer International Publishing, Cham
- Korczyński M, Nosyk Y, Lone Q, Skwarek M, Jonglez B, Duda A (2020b) Inferring the deployment of inbound source address validation using DNS resolvers. In: Proceedings of the Applied Networking Research Workshop, ANRW’20. ACM, pp 9–11
- Korczyński M, Nosyk Y, Lone Q, Skwarek M, Jonglez B, Duda A (2020c) The closed resolver project: measuring the deployment of source address validation of inbound traffic. In: CORR. <https://arxiv.org/abs/2006.05277>
- Kottler S (2018) February 28th DDoS incident report. <https://github.blog/2018-03-01-ddos-incident-report/>
- Kührer M, Hupperich T, Rossow C, Holz T (2014) Exit from hell? Reducing the impact of amplification DDoS attacks. In: USENIX Conference on Security Symposium
- Lichtblau F, Streibelt F, Krüger T, Richter P, Feldmann A (2017) Detection, classification, and analysis of inter-domain traffic with spoofed source IP addresses. In: Internet Measurement Conference, ACM
- Lone Q, Luckie M, Korczyński M, van Eeten M (2017) Using loops observed in traceroute to infer the ability to spoof. In: Passive and Active Measurement Conference. Springer International Publishing, Cham
- Lone Q, Luckie M, Korczyński M, Asghari H, Javed M, van Eeten M (2018) Using crowdsourcing marketplaces for network measurements: the case of spoofer. In: Traffic Monitoring and Analysis Conference
- Lone Q, Korczyński M, Gañán C, van Eeten M (2020) SAVing the internet: explaining the adoption of source address validation by internet service providers. In: Workshop on the Economics of Information Security
- Luckie M, Beverly R, Koga R, Keys K, Kroll J, claffy k (2019) Network hygiene, incentives, and regulation: deployment of source address validation in the internet. In: Computer and Communications Security Conference, ACM
- Luo X, Wang L, Xu Z, Chen K, Yang J, Tian T (2018) A large scale analysis of DNS water torture attack. In: Conference on Computer Science and Artificial Intelligence
- MANRS (2020) Mutually agreed norms for routing security. <https://www.manrs.org/>
- Mauch J (2013) Spoofing ASNs. <http://seclists.org/nanog/2013/Aug/132>

- Müller LF, Luckie MJ, Huffaker B, Claffy KC, Barcellos MP (2019) Challenges in inferring spoofed traffic at IXPs. In: Conference on Emerging Networking Experiments And Technologies, ACM
- Rossow C (2014) Amplification hell: revisiting network protocols for DDoS abuse. In: Network and Distributed System Security Symposium
- Senie D, Ferguson P (2000) Network ingress filtering: defeating denial of service attacks which employ IP source address spoofing. RFC 2827. <https://rfc-editor.org/rfc/rfc2827.txt>
- Shafir L, Afek Y, Bremler-Barr A (2020) NXNSAttack: recursive DNS inefficiencies and vulnerabilities. In: USENIX Security Symposium
- Spoofers Project (2020) The spoofers project. <https://www.caida.org/projects/spoofers/>
- Vixie P, Thomson S, Rekhter Y, Bound J (1997) Dynamic updates in the domain name system (DNS UPDATE). Internet RFC 2136