Inferring the Deployment of Inbound Source Address Validation Using DNS Resolvers

Maciej Korczyński, Yevheniya Nosyk, Qasim Lone, Marcin Skwarek, Baptiste Jonglez, Andrzej Duda

► To cite this version:

Maciej Korczyński, Yevheniya Nosyk, Qasim Lone, Marcin Skwarek, Baptiste Jonglez, et al.. Inferring the Deployment of Inbound Source Address Validation Using DNS Resolvers. ANRW 2020 - Applied Networking Research Workshop, Jul 2020, Madrid, Spain. pp.9-11, 10.1145/3404868.3406668 . hal-02894078

HAL Id: hal-02894078

<https://hal.science/hal-02894078>

Submitted on 8 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Inferring the Deployment of Inbound Source Address Validation Using DNS Resolvers

Extended Abstract*

Maciej Korczyński
Université Grenoble Alpes

Yevheniya Nosyk
Université Grenoble Alpes

Qasim Lone
Deft University of Technology

Marcin Skwarek
Université Grenoble Alpes

Baptiste Jonglez
Université Grenoble Alpes

Andrzej Duda
Université Grenoble Alpes

ABSTRACT

This paper reports on the first Internet-wide active measurement study to enumerate networks not filtering *incoming packets* based on their source address. Our method identifies closed and open DNS resolvers handling requests from the outside of the network with the source address in the prefix of the tested network. The study gives the most complete picture of the *inbound Source Address Validation* deployment at network providers: 32,673 IPv4 ASes and 197,641 IPv4 BGP prefixes are vulnerable to spoofing of inbound traffic.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Security**.

1 INTRODUCTION

The Internet relies on IP packets containing source and destination addresses in packet headers. However, there is no packet-level authentication mechanism to ensure that the source addresses are genuine [2]. The modification of a source IP address, referred to as “IP spoofing”, is leveraged in Distributed Denial-of-Service (DDoS) attacks, and in particular, the reflection attacks [1]. As we cannot prevent packet header modification, a means to block spoofed packets is packet filtering at the network edge, formalized in RFC 2827 and called *Source Address Validation* (SAV) [16].

The role of IP spoofing in cyberattacks drives the need to estimate the level of SAV deployment by network providers.

*Based on a previously published paper [7] at PAM 2020.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ANRW '20, July 27–30, 2020, Online (Meetecho), Spain

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8039-3/20/07...\$15.00

<https://doi.org/10.1145/3404868.3406668>

There exist methods aimed at enumerating networks without packet filtering [1–3, 8–15]. However, a great majority of the existing work concentrates on *outbound* SAV, the root of DDoS attacks [8]. While less obvious, the lack of *inbound* filtering enables an attacker to appear as an internal host of a network and may reveal valuable information about the network infrastructure. Inbound IP spoofing may serve as a vector for zone poisoning [5], cache poisoning [4], or recent NXNSAttack [17], even if the DNS server is correctly configured as a closed resolver.

In this paper, we report on the results of the Closed Resolver Project [6, 7, 19]. We propose a new method to identify networks not filtering inbound traffic based on source IP addresses. We perform an Internet-wide scan of the IPv4 address space to identify closed and open DNS resolvers in each routable network. We achieve this goal by sending DNS A requests with spoofed source IP addresses for which the destination is every host of every routing prefix and the source is the next host in the same network. We control the authoritative name server for the queried domains and observe from which networks it receives the requests. This method identifies networks not performing filtering of *incoming packets* without the need for a vantage point inside the network itself.

The above method when applied alone shows the absence of inbound SAV at the network edge. In parallel, we send subsequent unspoofed DNS A record requests to identify open resolvers at the scale of the Internet. If open resolvers reply to the unspoofed requests but not to the spoofed ones, we infer the presence of SAV for incoming traffic either at the network edge or in transit networks. By doing this, we detect both the absence and the presence of inbound packet filtering.

2 METHODOLOGY

2.1 Spoofing Scan

Figure 1 illustrates the idea of the proposed method in detail. We have developed an efficient scanner that sends handcrafted DNS A record request packets [18]. We run the scanner on a machine inside a network that does not deploy

Figure 1: Inbound spoofing scan setup.

outbound SAV so that we can send packets with spoofed IP addresses. We set up a DNS server authoritative for the `drakkardns.com` domain to capture the traffic related to our scans. When a resolver inside a network vulnerable to inbound spoofing performs query resolution, we observe it on our authoritative DNS server. To prevent caching and to be able to identify the true originator in case of forwarding, each time we query a unique subdomain composed of a random string, the hex-encoded resolver IP address (the destination of our query), a scan identifier, and the domain name itself, for example, `qGPDBe.02ae52c7.s1.drakkardns.com`.

Figure 1 shows the scanning setup for the `1.2.3.0/24` network. In step 1, the scanner sends one spoofed packet to each host of this network, thus packets to 254 destinations in total. The spoofed source IP address is always the next one after the destination. When the spoofed DNS packet has not been filtered anywhere in transit and there is no packet filtering at the edge, then nothing prevents it from entering the network. If the packet destination is `1.2.3.5`, the address of the local resolver (step 2), it receives a DNS A record request from what looks to be another host on the same network and performs query resolution. If the destination is not the local resolver, it will drop the packet. However, the scanner will eventually reach all the hosts on the network and the local resolver if there is one.

In this study, we distinguish between two types of local resolvers: forwarders (or proxies) that forward queries to other recursive resolvers and non-forwarders that recursively resolve queries they receive. Therefore, the non-forwarding local resolver (`1.2.3.5`) inspects the query that looks as if it was sent from `1.2.3.6` and performs the resolution by iteratively querying the root (step 3) and the top-level domain name (step 4) servers until it reaches our authoritative DNS server in step 5. Alternatively, it forwards the query to another recursive resolver that repeats the same procedure as described above for non-forwarders. In step 6, the DNS A

query response is sent to the spoofed source (`1.2.3.6`).

2.2 Open Resolver Scan

In parallel, we perform an open resolver scan by sending DNS A requests with the genuine source IP address of the scanner. To avoid temporal changes, we send a non-spoofed query just after the spoofed one to the same host. The format of a non-spoofed query is almost identical to the spoofed one. The only difference is the scan identifier, for example, `qGPDBe.02ae52c7.n1.drakkardns.com`. If we receive a request on our authoritative DNS server, it means that we have reached an open resolver. Moreover, if this open resolver did not resolve a spoofed query, we infer the presence of inbound SAV either in transit or at the tested network edge.

3 RESULTS AND CONCLUSIONS

Our scans have revealed 4,589,251 closed DNS resolvers, 4,213,192 of which are forwarders. We have also identified 4,022,711 open resolvers. We found that 32,673 ASes (49%), 197,641 BGP prefixes (23%), and 959,666 /24 IPv4 networks are fully or partially vulnerable to spoofing of inbound traffic. We have compared the results of spoofing and open resolver scans to reveal the absence and the *presence* of inbound SAV. We found that 19,090 of ASes are not vulnerable to inbound spoofing. However, 14,382 (38.47%) of all tested autonomous systems show inconsistent results possibly due to different filtering policies at upstream providers for multi-homed customer ASes, or measurement errors.

We have retrieved the Spoofer data and deployed a method proposed by Mauch [14] to infer the absence and the presence of *outbound* SAV. In this way, we studied the policies of the SAV deployment per provider in both directions and concluded that inbound filtering is less deployed than outbound.

Previous work demonstrated the difficulty in incentivizing providers to deploy filtering for outbound traffic due to misaligned economic incentives: implementing SAV for outbound traffic benefits other networks and not the network of the deployment [13]. This work shows how the deployment of SAV for inbound traffic protects the provider network.

We have started longitudinal measurements to infer the deployment of SAV in both IPv4 and IPv6 address spaces [6] and plan to notify all parties affected by the vulnerability.

ACKNOWLEDGEMENTS

This work has been carried out in the framework of the PrevDDoS project funded by the IDEX Université Grenoble Alpes "Initiative de Recherche Scientifique (IRS)" and partially supported by the Grenoble Alpes Cybersecurity Institute CYBER@ALPS under contract ANR-15-IDEX-02, PERSYVAL-Lab under contract ANR-11-LABX-0025-01, and DiNS under contract ANR-19-CE25-0009-01.

REFERENCES

- [1] Robert Beverly and Steven Bauer. 2005. The Spoofer Project: Inferring the Extent of Source Address Filtering on the Internet. In *USENIX Steps to Reducing Unwanted Traffic on the Internet (SRUTI) Workshop*.
- [2] R. Beverly, A. Berger, Y. Hyun, and k. claffy. 2009. Understanding the Efficacy of Deployed Internet Source Address Validation Filtering. In *Internet Measurement Conference*. ACM.
- [3] CAIDA. 2020. *The Spoofer Project*. <https://www.caida.org/projects/spoofer/>
- [4] Dan Kaminsky. 2008. It's the End of the Cache as We Know It. <https://www.slideshare.net/dakami/dmk-bo2-k8>.
- [5] Maciej Korczyński, Michał Król, and Michel van Eeten. 2016. Zone Poisoning: The How and Where of Non-Secure DNS Dynamic Updates. In *Internet Measurement Conference*. ACM.
- [6] Maciej Korczyński, Yevheniya Nosyk, Qasim Lone, Marcin Skwarek, Baptiste Jonglez, and Andrzej Duda. 2020. The Closed Resolver Project: Measuring the Deployment of Source Address Validation of Inbound Traffic. *arXiv:2006.05277 [cs.NI]*
- [7] Maciej Korczyński, Yevheniya Nosyk, Qasim Lone, Marcin Skwarek, Baptiste Jonglez, and Andrzej Duda. 2020. Don't Forget to Lock the Front Door! Inferring the Deployment of Source Address Validation of Inbound Traffic. In *Passive and Active Measurement Conference*. Springer, 107–121.
- [8] Marc Kührer, Thomas Hupperich, Christian Rossow, and Thorsten Holz. 2014. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *USENIX Conference on Security Symposium*.
- [9] Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Philipp Richter, and Anja Feldmann. 2017. Detection, Classification, and Analysis of Inter-domain Traffic with Spoofed Source IP Addresses. In *Internet Measurement Conference*. ACM.
- [10] Qasim Lone, Maciej Korczyński, Carlos Gañán, and Michel van Eeten. 2020. SAVING the Internet: Explaining the Adoption of Source Address Validation by Internet Service Providers. In *Workshop on the Economics of Information Security*.
- [11] Qasim Lone, Matthew Luckie, Maciej Korczyński, Hadi Asghari, Mobin Javed, and Michel van Eeten. 2018. Using Crowdsourcing Marketplaces for Network Measurements: The Case of Spoofer. In *Traffic Monitoring and Analysis Conference*.
- [12] Qasim Lone, Matthew Luckie, Maciej Korczyński, and Michel van Eeten. 2017. Using Loops Observed in Traceroute to Infer the Ability to Spoof. In *Passive and Active Measurement Conference*. Springer.
- [13] M. Luckie, R. Beverly, R. Koga, K. Keys, J. Kroll, and k claffy. 2019. Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In *Computer and Communications Security Conference (CCS)*. ACM.
- [14] Jared Mauch. 2013. Spoofing ASNs. <http://seclists.org/nanog/2013/Aug/132>.
- [15] Lucas F. Müller, Matthew J. Luckie, Bradley Huffaker, kc claffy, and Marinho P. Barcellos. 2019. Challenges in Inferring Spoofed Traffic at IXPs. In *Conference on Emerging Networking Experiments And Technologies (CoNEXT)*. ACM, 96–109.
- [16] Daniel Senie and Paul Ferguson. 2000. Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing. RFC 2827. <https://rfc-editor.org/rfc/rfc2827.txt>
- [17] Lior Shafir, Yehuda Afek, and Anat Bremner-Barr. 2020. NXNSAttack: Recursive DNS Inefficiencies and Vulnerabilities. In *USENIX Security Symposium*.
- [18] Marcin Skwarek, Maciej Korczyński, Wojciech Mazurczyk, and Andrzej Duda. 2019. Characterizing Vulnerability of DNS AXFR Transfers with Global-Scale Scanning. In *2019 IEEE Security and Privacy Workshops*. 193–198.
- [19] The Closed Resolver Project. 2020. <https://closedresolver.com>.