Study on Anti-Spoofing in the Region

Author: Paula Venosa, Nicolás Macia, Guillermo Pereyra
Coordination and Revision: Graciela Martínez
Edition: Communications Department
Department: LACNIC's CSIRT Department

December 2022

lacnic
www.lacnic.net

Executive Summary

This study managed to identify sources of information to assess the proper implementation of good practices to mitigate IP spoofing in networks across Latin America and the Caribbean. Based on the analysis of the data we obtained, this document recommends actions that will allow organizations to detect and mitigate this problem. Thus, the report does not only show the status of adherence to good practices, but it also shows that it is possible to improve IP address spoofing prevention in the region.

These statistics provide LACNIC CSIRT with information to work directly with member organizations and reduce the risk of the resources assigned to them being used for malicious purposes.

Abstract

This report is intended as a single source of information to understand IP spoofing and its related issues. It is a starting point to raise awareness on this topic, both among LACNIC members as well as throughout the community.

The report presents the results of the survey conducted to determine the existence of networks that enable IP spoofing, and which can be used to launch attacks in Latin America and the Caribbean. It describes the most common attacks, in addition to the methodology that was applied, the sources that were examined, the results that were obtained after an initial survey, the most relevant recommendations for their correction, and the conclusions of the study.

It also presents the community with global data on compliance with good practices to mitigate IP spoofing in the region, based on data obtained from the CAIDA Spoofer and the Closed Resolver Project.

Goal

To contribute to the enhancement of security in the LACNIC region, encouraging member organizations to implement good security practices that will contribute to Internet stability and resilience.

Motivation

To improve the cybersecurity levels of the IP resources under LACNIC's administration.

To minimize the potential use of IP spoofing techniques to launch attacks on both the infrastructure that underpins the Internet as well as on our members' networks, by raising awareness and providing tools that the LACNIC community can use to address this issue.

IP Spoofing

What is IP spoofing?

IP address spoofing or IP spoofing is masquerading as another IP address. It refers to the creation of IP packets with a false source IP address, for the purpose of impersonating another computing system. IP spoofing allows malicious actors to launch different types of attacks and makes their detection very difficult.

In IP spoofing, a malicious actor uses tools that modify the IP address of the sender of the packets so that to the recipient it appears that the packets came from another source.

IP spoofing is possible when proper filtering policies are not implemented to prevent this type of attacks.[1] [2] These filtering policies should check the source address of the packets. To refer to these policies, we will use the term source address validation (SAV) or SAV filters.

An organization may unknowingly allow IP spoofing, both for incoming traffic (inbound IP spoofing) as well as for traffic bound for the Internet (outbound IP spoofing).

Inbound IP spoofing

The possibility that traffic from external autonomous systems will enter an organization by appearing to have originated from IP addresses within the organization itself or in bogon networks.[3] This should not occur and such traffic should be filtered, as:

---

[1] RFC 2827 - Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing ‹https://www.rfc-editor.org/rfc/rfc2827›

[2] RFC 3704 - Ingress Filtering for Multihomed Networks ‹https://www.rfc-editor.org/rfc/rfc3704.html›

[3] Bogon Networks - ‹https://team-cymru.com/community-services/bogon-reference/›

- The organization's IP addresses should be the destination and not the source of incoming traffic. When it uses the organization's own IP addresses, inbound IP spoofing originates various attacks that may affect its users and services.
- Bogon networks should never appear in internet routing tables.

Inbound IP spoofing can occur when an organization does not deploy SAV for inbound traffic (inbound SAV).

Outbound IP spoofing

The possibility that IP traffic with source IP addresses that do not belong to the organization itself may leave the organization towards external autonomous systems. These situations should never occur, and these packets should be filtered, as they allow various attacks to be launched that may affect the availability of third-party services.

Outbound IP can occur when an organization does not deploy SAV for outbound traffic (outbound SAV).

Which protocols are leveraged for IP spoofing?

IP spoofing takes place in the network layer, at the IP address level. Depending on the attack, however, a different protocol may be encapsulated in the IP packets.

For example, flooding DoS[4] attacks can use various protocols:

- HTTP to saturate the service.
- TCP with specific flags, such as syn, fin, first, etc., to saturate the service.
- ICMP to increase the volume of the attack and saturate the capacity of the network or server.
- UDP to increase the volume of the attack and saturate the capacity of the network or server.

Distributed denial-of-service (DDoS) and other types of attacks that need to interact with services using spoofed traffic use application protocols transported over UDP.

Applications that use these protocols provide what are known as reflective services. These services use UDP as the transport protocol and do not require handshaking prior to exchanging data.

---

[4] DoS / DDoS - ‹https://www.osi.es/es/actualidad/blog/2018/08/21/que-son-los-ataques-dos-y-ddos›

What type of attacks leverage IP spoofing and how do they do it?

Flooding DoS attacks

A DoS attack seeks to compromise the availability of a service in the target organization. If a malicious actor controls a host in an organization that does not deploy outbound SAV, they may hide the source of the attack by performing IP spoofing.

Amplification DDoS attacks

For DDoS attacks to be possible, two conditions must be met. First, the organization from which the attack is launched must not have any outbound SAV in place; second, vulnerable reflective and amplifiable services that can be accessed from the Internet must be identified.[5]

The report titled "Study of Internet Resources in the LACNIC Service Region that Use UDP Protocols and Enable DDoS Attacks" addresses the latter issue.[6] It describes the vulnerabilities that can enable distributed denial-of-service (DDoS) attacks.

In short, the malicious actor must:

1. Identify vulnerable reflective and scalable services that can be accessed from the Internet.
2. Send spoofed requests to previously identified vulnerable services (present on the Internet due to a lack of knowledge), using the IP address of the attack target as the source IP address.
3. Perform the previous step in a coordinated manner from multiple hosts, in organizations that do not apply outbound SAV, thus making it possible to increase the volume of the attack proportionally to the number of hosts used by the malicious actor.

---

[5] ‹https://www.cisa.gov/uscert/ncas/alerts/TA14-017A›

[6] ‹https://imasd.lacnic.net/reportes/ciberseguridad/estudio-de-recursos-de-internet-presentes-en-la-region-de-lacnic-con-protocolos-udp-que-posibilitan-ataques-de-tipo-ddos.pdf ›

NSNXAttacks

Unlike the attacks described above, an NSNXAttack takes advantage of the fact that the victim organization does not deploy inbound SAV. In these attacks, a malicious actor sends queries to the DNS resolver as if they were using a host belonging to the organization. This type of attack specifically seeks to affect the availability of the service provided by the organization's recursive DNS.[7]

DNS cache poisoning

DNS cache poisoning[8] is the practice of providing a domain name server (DNS) with data that does not originate from authoritative sources. DNS cache poisoning can occur regardless of whether the victim organization deploys inbound SAV. However, failure to deploy inbound SAV may facilitate an attack.

DNS cache poisoning affects the integrity of the data provided by the DNS service, creating what can be the first element of more complex attacks, for example, phishing, bank fraud,[9] or website defacement attacks.[10]

Other possible attacks

Because inbound IP spoofing allows a malicious actor to interact with internal services using local addresses, other application protocols may also allow various attack vectors to leverage the fact that the zero-trust paradigm is not applied. [11] [12] [13] [14]

Benefits of implementing actions against IP spoofing

LACNIC members can benefit from the implementation of actions to prevent IP spoofing by mitigating various attacks that take advantage of inbound IP spoofing against their organization. Attacks such as NSNXAttacks and DNS cache poisoning, as well as others that may be devised in the future, are examples of attacks that take advantage of the trust

---

7 ‹https://www.nxnsattack.com/shafir2020-nxnsattack-paper.pdf›

8 ‹https://es.wikipedia.org/wiki/Envenenamiento_de_DNS›

9 ‹https://www.zscaler.it/blogs/security-research/brazilian-bank-targeted-phishing-site-and-dns-poisoning›

10 ‹https://www.theguardian.com/technology/2017/aug/31/wikileaks-hacked-ourmine-group-julian-assange-dns-attack›

11 ‹https://storage.googleapis.com/pub-tools-public-publication-data/pdf/43231.pdf›

12 ‹https://cloud.google.com/blog/topics/developers-practitioners/what-zero-trust-identity-security›

13 ‹https://www.cloudflare.com/es-es/learning/security/glossary/what-is-zero-trust/›

14 ‹https://www.akamai.com/es/our-thinking/zero-trust/zero-trust-security-model›

placed on internal communications within environments that are not in line with the zero-trust paradigm.

Likewise, the Internet community as a whole can benefit if the different actors take actions to prevent IP spoofing, limiting the networks from which DDoS attacks can be launched.

Available sources of information

For this study, active data sources were used that could provide relevant information on the status of IP spoofing in the region.

These sources were:

- ‹https://www.caida.org/projects/spoofer/›
- ‹https://closedresolver.korlabs.io/›

The CAIDA Spoofer Project

The CAIDA Spoofer[15] Project is a mandatory reference for anyone looking to obtain data on issues related to IP spoofing. It has been running since 2006 and has a client that can be used to perform tests. These tests allow assessing whether a network allows both outbound IP spoofing and inbound IP spoofing.

The following image shows the architecture of the CAIDA Spoofer, its components, and their interaction:

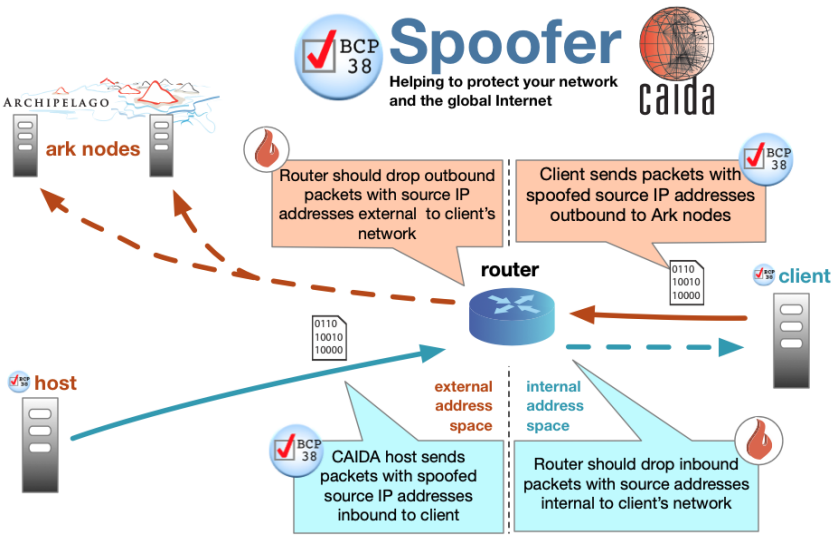

Figure 1- CAIDA Spoofer ‹https://www.caida.org/projects/spoofer/images/spoofer-diagramer.png›

The client software performs the tests to assess the organization's filtering policies. The client sends packets with spoofed source IP addresses outbound to the ark nodes in the CAIDA archipelago.

15 ‹https://www.caida.org/projects/spoofer/›

If filtering policies are applied, the organization's router drops outbound packets with IP addresses external to the client's network.

Ark nodes act like honeypots. If they receive packets that originate in the client with spoofed IP addresses, it can be concluded that the organization does not filter outbound IP spoofing. If the ark nodes do not receive these packets, then the organization is filtering outbound IP spoofing.

The CAIDA host is the external device from which the inbound IP spoofing tests are performed. From there, packets with spoofed source IP addresses are sent to the client.

When filtering policies are applied, the organization's router drops inbound packets with source IP addresses that are not recognized as internal to the client's network.

For each country, the project provides information on the number of IP blocks that have been probed, the number of blocks that are not filtering outbound IP spoofing, and the number of blocks that are. It also provides reports containing evidence of remediation at country level.

The advantages of this project include:

- It provides the user with information on the existence of inbound and outbound SAV in the organization.
- It allows assessing IP spoofing, both for IPv4 and IPv6 networks.
- The client application installed on the device automatically probes the existence of inbound and outbound SAV in the networks to which the device connects.

Its disadvantages include:

- Public data shows only outbound IP spoofing.
- It requires the installation of a client software.
- It contains partial information for IPv4 and IPv6 networks.

The Closed Resolver Project

In 2019, The Closed Resolver Project[16] scanned the entire IPv4 address space using DNS to assess the presence of proper filters against inbound IP spoofing. DNS was used to send queries that appeared to have originated in hosts within the same organization where the DNS server was running. When the recipient was a recursive DNS server, detecting

16 ‹https://closedresolver.korlabs.io/›

inbound IP spoofing was possible using the sequence illustrated in the image below (taken from the article that describes the goal of the project).[17]

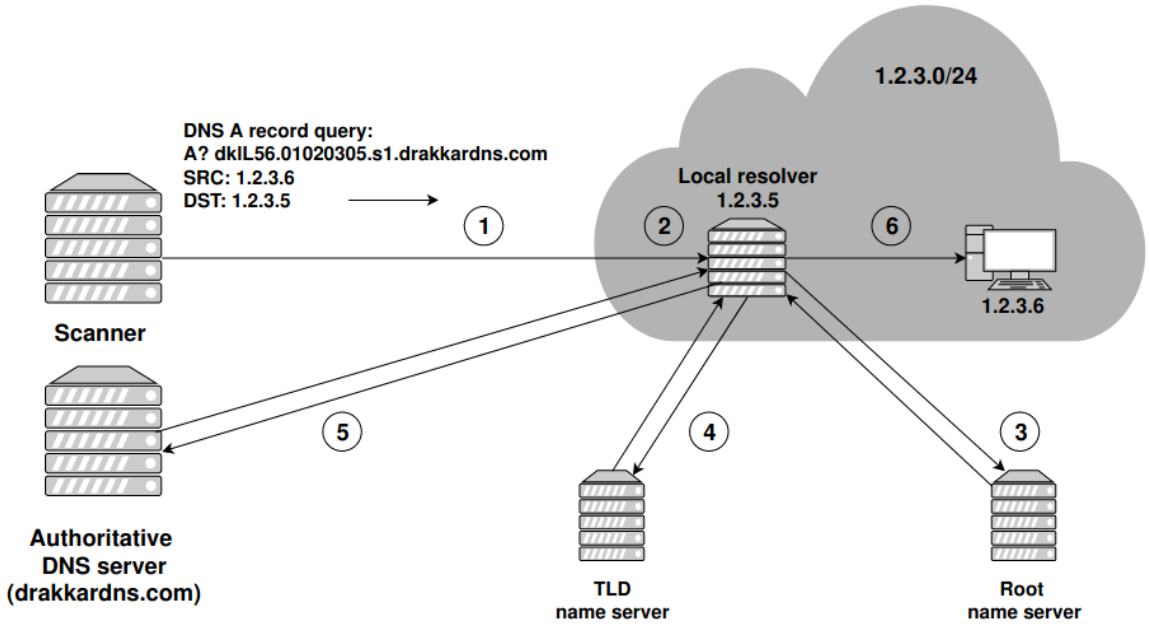

Fig. 2- Methodology used by The Closed Resolver Project to assess inbound IP spoofing ‹https://arxiv.org/pdf/2002.00441.pdf›

The methodology shown in Figure 2 is as follows:

1. DNS requests are sent to different IPv4 addresses (1.2.3.5 in the example) from IP addresses within the organization itself (1.2.3.6 in the example).
2. If inbound SAV is not implemented, inbound IP spoofing will be possible and the packet will reach its destination (1.2.3.5 in the example).
3. If there is a DNS resolver at the DNS request's destination address, it will try to resolve the query (in the example, the A record of dkIL56.01020305.s1.drakkardns.com is queried). To do so, the resolver will initially contact a root server that will provide information on the name servers (NS) associated with the top-level domain (TLD).
4. Next, one of the name servers associated with the corresponding TLD will be contacted. The TLD's response will provide information on the NS's associated with the drakkardns.com domain.
5. Because the drakkardns.com domain in the example is controlled by The Closed Resolver Project, upon receiving the query it is possible to confirm that it originated in a network that does has not implemented inbound SAV.

---

17 Figure taken from ‹https://arxiv.org/pdf/2002.00441.pdf›

6. Finally, the DNS server provides the requested information to whoever it thinks sent the query (1.2.3.6 in the example).

The advantages of this project include:

- It contains complete data for IPv4 networks.
- It does not require the installation of any software.

Its disadvantages include:

- It only provides information on good practices related to the application of inbound SAV.
- Only partial data is available for IPv6. Because it is not possible to scan the entire address space, the hitlist[18] service is used to test the IPv6 addresses that respond.

How to determine whether your organization implements SAV

To quickly determine whether an organization is properly implementing inbound and outbound SAV, a manual assessment may be performed. This would include checking whether the following situations are possible:

- Inbound IP spoofing: sending traffic with source IP addresses that are part of the IP address blocks assigned to the organization from the Internet to the organization.
- Outbound IP spoofing: sending traffic with source IP addresses that are part of the IP address blocks assigned to the organization from the organization to the Internet.

Manual IP spoofing detection mechanism

Manually assessing whether IP spoofing is possible is a convenient way to understand and check the problem. To manually detect if an organization allows both inbound and outbound IP spoofing, we need:

- A tool for generating custom traffic. In this case, Hping3 will be used.[19]
- A tool to monitor the arrival of the custom traffic. In this case, Tcpdump will be used.[20]

---

18 ‹https://ipv6hitlist.github.io/›

19 Hping3 ‹http://www.hping.org/›

20 Tcpdump ‹https://www.tcpdump.org/›

- A network device connected to the network of the organization to be evaluated and which can receive unfiltered traffic to a specific service (for example: TCP/443).
- A network device connected to an autonomous system that is external to the organization to be evaluated.

Manual inbound IP spoofing check

The lack of inbound SAV enables inbound IP spoofing. This can be manually verified by generating spoofed traffic on a device connected to an external autonomous system using the organization's own IP addresses as source IP addresses.

The traffic must be directed to a device connected to the organization's network. If the custom traffic reaches this last device, we can conclude that inbound IP spoofing is possible.

The steps to perform the manual verification are as follows:

Device connected to an external autonomous system

```
# Assuming that
# the IP of the device connected to an external autonomous system is y.y.y.y
# the IP of the device in the organization is x.x.x.1
# the IP that will be used as the source IP for the test is x.x.x.2
# Initial verification test: to check that the traffic arrives
hping -S -p 443 -c 1 x.x.x.1
# Inbound IP spoofing test
hping -S -p 443 -c 1 -a x.x.x.2 x.x.x.1
```

Device connected to the organization to be assessed

```
# Assuming that the network interface is eth0 and using filters to observe
# only the traffic of interest
tcpdump -n -i eth0 host y.y.y.y or host x.x.x.2
```

Depending on the activity shown by the **tcpdump** command, we may be able to conclude whether inbound IP spoofing is possible. Based on these observations, the following scenarios are possible:

- Inbound IP spoofing is possible: This can be concluded if the tcpdump command returns the traffic associated with the initial verification test as well as the traffic associated with the inbound IP spoofing test.

```
IP y.y.y.y.7323 > x.x.x.1.443: Flags [S], seq 45611324, length 0
```

```
      IP x.x.x.1.443 > y.y.y.y.7323: Flags [R.], seq 0, ack 45611325, length
0

      IP x.x.x.2.1899 > x.x.x.1.443: Flags [S], seq 883275884, length 0

      IP x.x.x.1.443 > x.x.x.2.1899: Flags [R.], seq 0, ack 883275885,
length 0
```

- Inbound IP spoofing is not possible: This can be concluded if the tcpdump command returns only the traffic associated with the initial verification test.

```
      IP y.y.y.y.7323 > x.x.x.1.443: Flags [S], seq 45611324, length 0

      IP x.x.x.1.443 > y.y.y.y.7323: Flags [R.], seq 0, ack 45611325, length
0
```

- Impossible to determine: If the command does not return the traffic associated with the initial verification test, we can conclude that a network filter is preventing us from performing the verification on device x.x.x.1, using the specified port (in the example, port TCP/443).

Manual outbound IP spoofing check

Similarly, the lack of outbound SAV enables outbound IP spoofing. This can be verified manually by generating spoofed traffic on a device connected to the organization that is being assessed, using IP addresses that do not belong to the organization as source IP addresses.

The traffic must be directed to a device connected to an external autonomous system. If the custom traffic reaches this last device, we can conclude that outbound IP spoofing is possible.

The steps for the manual verification are as follows:

Device connected to the organization to be assessed

```
# Assuming that
# the IP of the device in the organization is x.x.x.1
# the IP of the device connected to an external autonomous system is y.y.y.y
# the IP that will be used as the source IP for the test is z.z.z.z
# Initial verification test: to check that the traffic arrives
hping -S -p 443 -c 1 y.y.y.y
# Outbound IP spoofing test
hping -S -p 443 -c 1 -a z.z.z.z y.y.y.y
```

Device connected to an external autonomous system

```
# Assuming that the network interface is eth0 and using filters to observe
```

```
# only the traffic of interest
tcpdump -n -i eth0 host x.x.x.1 or host z.z.z.z
```

Depending on the activity shown by the **tcpdump** command, we may be able to conclude whether outbound IP spoofing is possible. Based on these observations, the following scenarios are possible:

- Outbound IP spoofing is possible: This can be concluded if the tcpdump command returns the traffic associated with the initial verification test as well as the traffic associated with the outbound IP spoofing test.

```
IP x.x.x.1.7323 > y.y.y.y.443: Flags [S], seq 45611324, length 0
IP y.y.y.y.443 > x.x.x.1.7323: Flags [R.], seq 0, ack 45611325, length
0
IP z.z.z.z.1899 > y.y.y.y.443: Flags [S], seq 883275884, length 0
IP y.y.y.y.443 > z.z.z.z.1899: Flags [R.], seq 0, ack 883275885,
length 0
```

- Outbound IP spoofing is not possible: This can be concluded if the tcpdump command returns only the traffic associated with the initial verification test.

```
IP x.x.x.1.7323 > y.y.y.y.443: Flags [S], seq 45611324, length 0
IP y.y.y.y.443 > x.x.x.1.7323: Flags [R.], seq 0, ack 45611325, length
0
```

- Impossible to determine: If the command does not return the traffic associated with the initial verification test, we can conclude that a network filter is preventing us from performing the verification.

Automatic IP spoofing detection mechanism

A manual assessment of IP spoofing is a practical way to understand the problem with the aim of mitigating it. However, the use of tools that automate IP spoofing detection is recommended for a full, in-depth understanding of the problem. There are currently two tools that help to detect IP spoofing automatically.

The Closed Resolver Project

The Closed Resolver Project is a web portal that provides information on networks that are vulnerable to inbound IP spoofing. The visitor's IP address is used to report the results.

Spoofer

This tool is provided by the CAIDA Spoofer Project.[21] This project is the main source of data on IP spoofing. The tool provided by the project allows obtaining the data that is offered to the community.

To perform the tests, the spoofer program must be installed.[22] The tests take into account a large number of particular situations to assess both inbound and outbound IP spoofing

Spoofer assesses the ability to perform IP spoofing with both IPv4 and IPv6 traffic. It also assesses whether it is possible to use IP addresses that are part of non-routable IP blocks (e.g., private IP blocks[23] in IPv4 or ULA[24] in IPv6) as source IP addresses.

While it also has a graphical interface, the tool performs periodic tests, thus allowing the assessment of the different networks to which the user connects.

Figure 3 shows the CAIDA Spoofer graphical user interface, which includes an option for users to start new tests, as well as a summary of the tests that have been performed and the option to access the full report.

---

21 ‹https://www.caida.org/projects/spoofer/#download-client-software›

22 ‹https://www.caida.org/projects/spoofer/#download-client-software›

23 ‹http://www.ietf.org/rfc/rfc1918.txt›

24 ‹https://www.ietf.org/rfc/rfc4193.txt

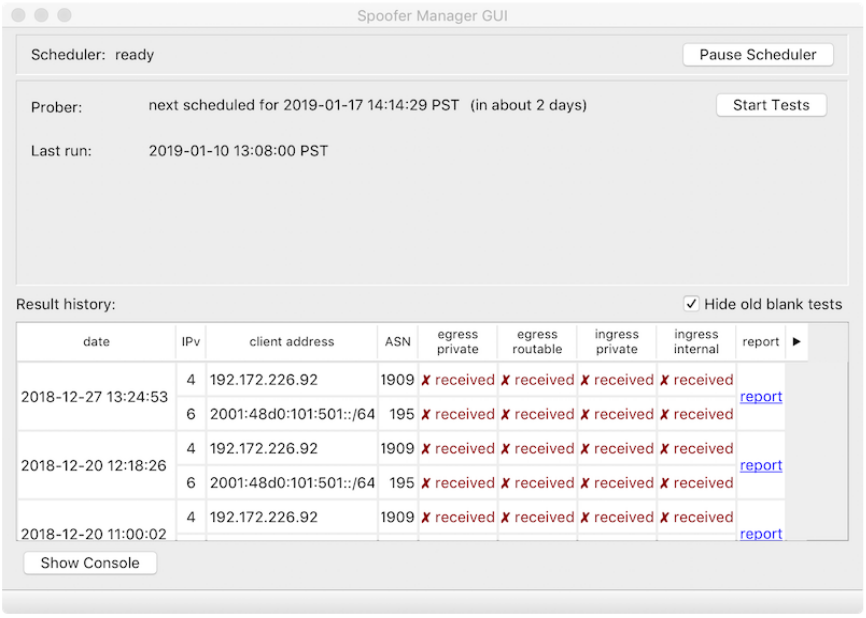

Fig. 3 - CAIDA Spoofer client GUI

The reports contain details on inbound and outbound IP spoofing. Figure 4 shows a summary of the test results for IPv4, IPv6, use of non-routable addresses, use of routable addresses and whether inbound and outbound IP spoofing are possible. Figure 5 shows the series of autonomous systems through which the spoofed packets were routed in the tests that were carried out.

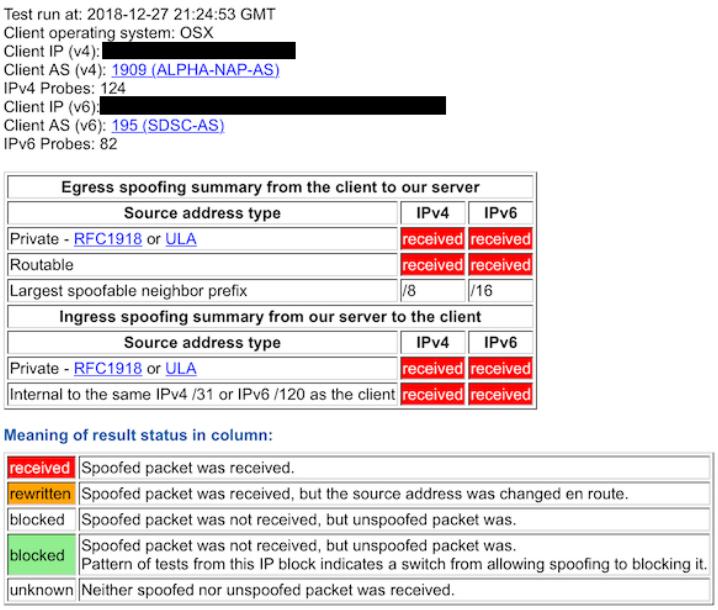

Test run at: 2018-12-27 21:24:53 GMT
Client operating system: OSX
Client IP (v4): [redacted]
Client AS (v4): 1909 (ALPHA-NAP-AS)
IPv4 Probes: 124
Client IP (v6): [redacted]
Client AS (v6): 195 (SDSC-AS)
IPv6 Probes: 82

Egress spoofing summary from the client to our server		
Source address type	**IPv4**	**IPv6**
Private - RFC1918 or ULA	received	received
Routable	received	received
Largest spoofable neighbor prefix	/8	/16
Ingress spoofing summary from our server to the client		
Source address type	**IPv4**	**IPv6**
Private - RFC1918 or ULA	received	received
Internal to the same IPv4 /31 or IPv6 /120 as the client	received	received

Meaning of result status in column:

received	Spoofed packet was received.
rewritten	Spoofed packet was received, but the source address was changed en route.
blocked	Spoofed packet was not received, but unspoofed packet was.
blocked	Spoofed packet was not received, but unspoofed packet was. Pattern of tests from this IP block indicates a switch from allowing spoofing to blocking it.
unknown	Neither spoofed nor unspoofed packet was received.

Fig. 4- Summary section of the CAIDA Spoofer client report

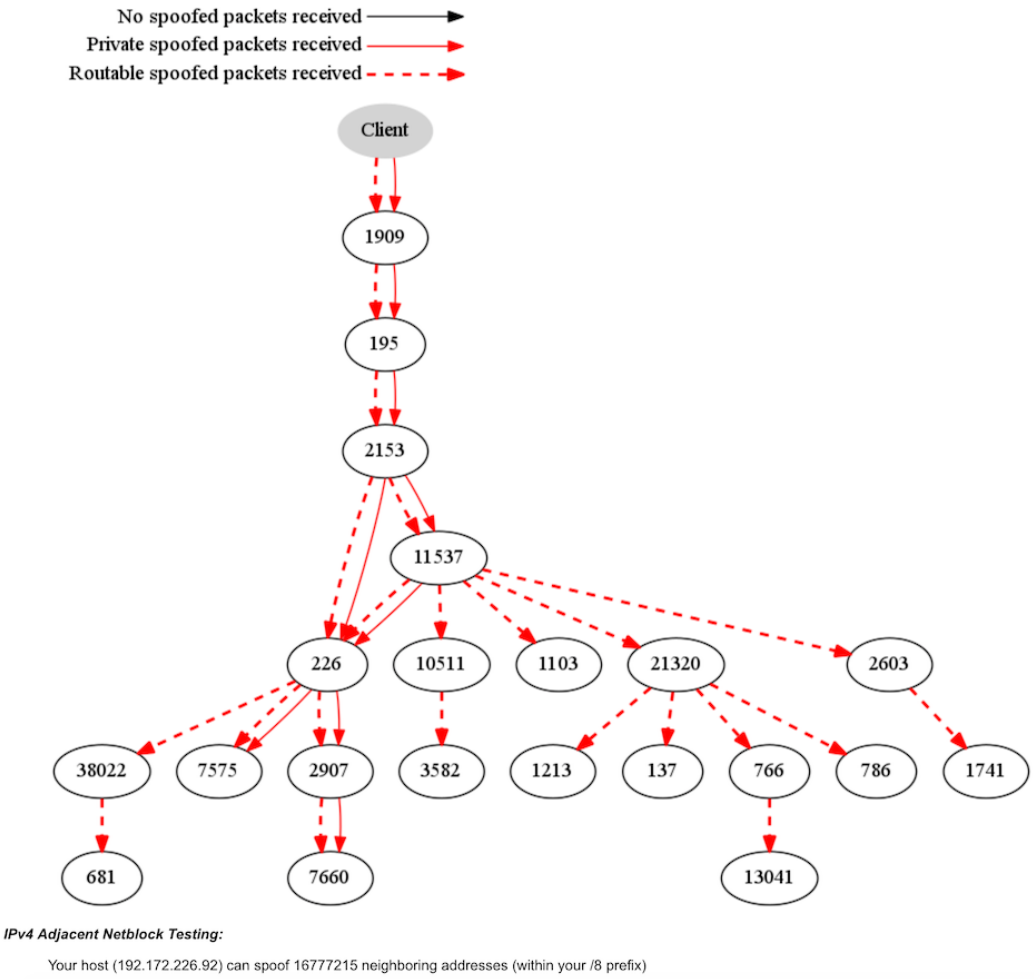

Fig. 5- IPv4 AS route section of the CAIDA Spoofer client report

Analysis of the implementation of anti-spoofing policies in the LAC region

This section describes the status of the region with respect to IP spoofing in networks involving resources under LACNIC administration.

It describes the main findings based on the information sources included in the section titled “Available sources of information.”

Outbound SAV implementation

The data in this section were obtained from ‹https://spoofer.caida.org/country_stats.php› on 17 October 2022

Observations regarding the tests conducted in the region

A total of 3082 blocks were evaluated in the region. For each AS, several IP blocks may have been assessed.

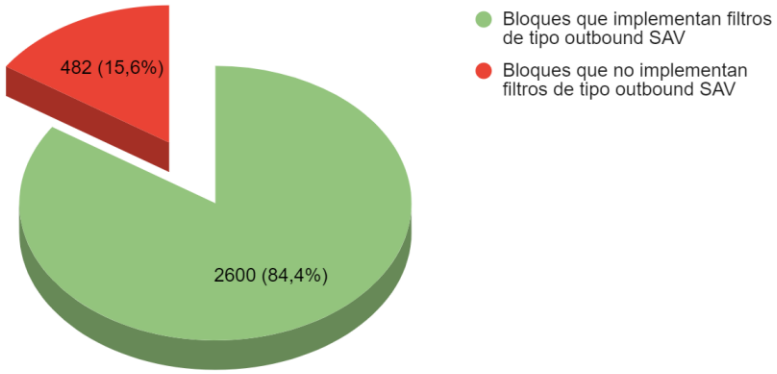

Figura 6

Figure 6 - Out of the total number of blocks that were evaluated, 84.4% implement outbound SAV, while the remaining 15.6% do not.

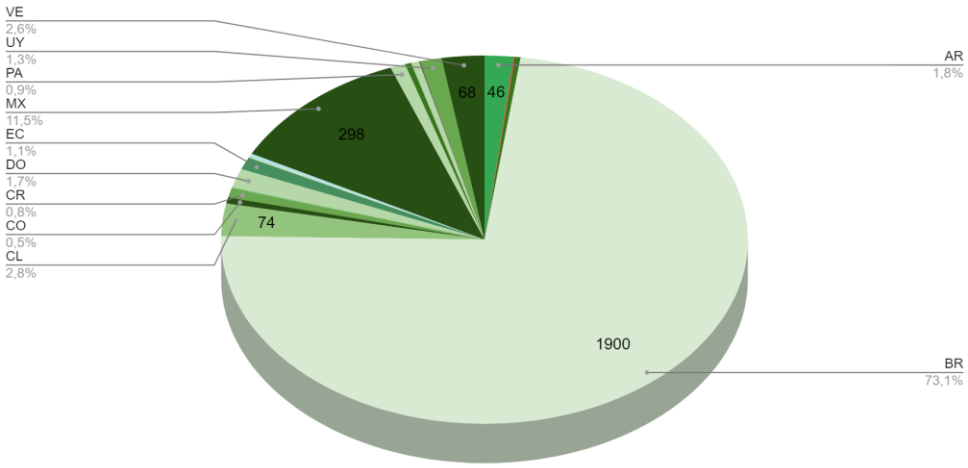

Figura 7

Figure 7 - Percentage of blocks that implement outbound SAV per country (out of the total of 84.4% shown in Figure 6).

The images above provide a first insight into the outbound SAV deployment in the region.

Observations regarding the distribution of the tests among the countries of the region

To assess the situation by country, it is important to consider the distribution of the blocks that were analyzed, i.e., the number of tests per country. Figure 8 shows the number of countries in the region where tests were and were not carried out. Figure 9 breaks down the number of tests by country.

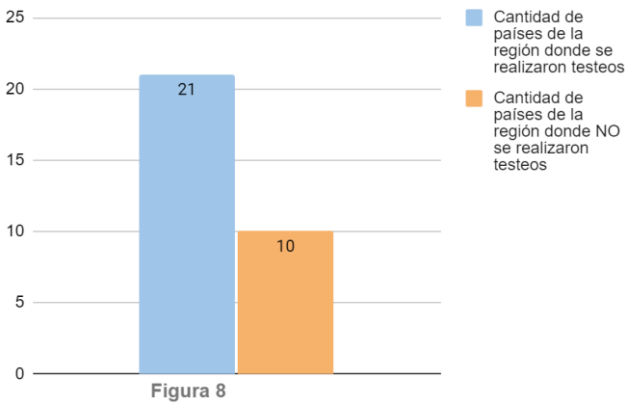

Figura 8

As Figure 8 shows, out of a total of 31 countries, tests to detect outbound SAV were carried out in 21. No tests were conducted in the remaining 10 countries.

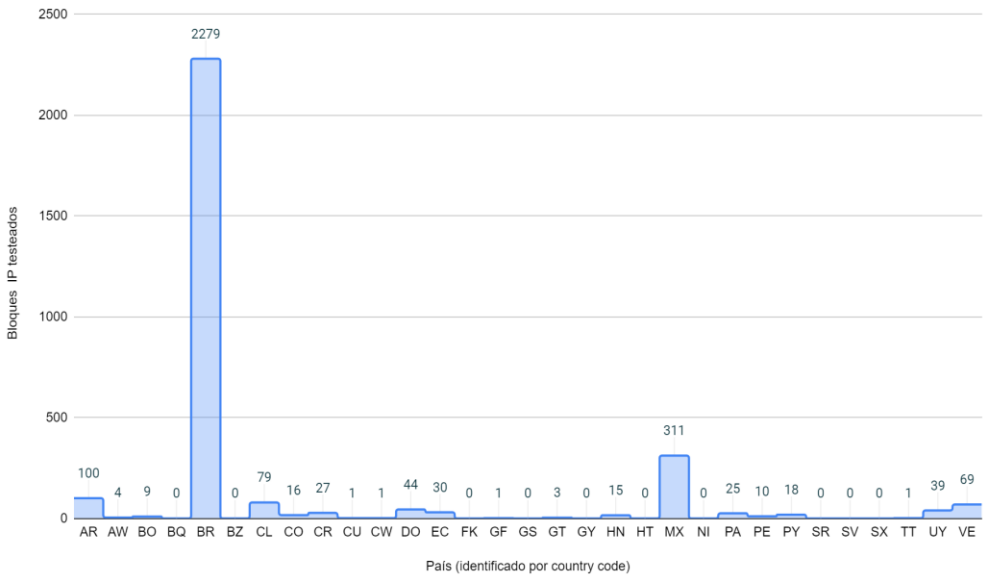

Figura 9

Figure 9 shows that Brazil was the country where the highest number of tests were performed (2279 IP blocks), followed by Mexico (311 IP blocks), and Argentina (100 IP blocks). It also shows the 10 countries where no tests were carried out: Bonaire, Saint Eustatius and Saba, Belize, Falkland Islands (Malvinas Islands), Guyana, South Georgia and the South Sandwich Islands, Haiti, Nicaragua, Suriname, El Salvador, and Saint Martin.

Observations regarding the implementation of outbound SAV in the blocks tested in each country

Figures 10, 11, 12, and 13 show the situation in each country. Countries where no testing was conducted are not included in these figures [Bonaire, Saint Eustatius and Saba, Belize, Falkland Islands (Islas Malvinas), Guyana, South Georgia and the South Sandwich Islands, Haiti, Nicaragua, Suriname, El Salvador, South Georgia Islands, and Saint Martin].

Each of these figures shows the number of IP blocks that do and do not implement outbound SAV out of the total number of blocks tested in each country:

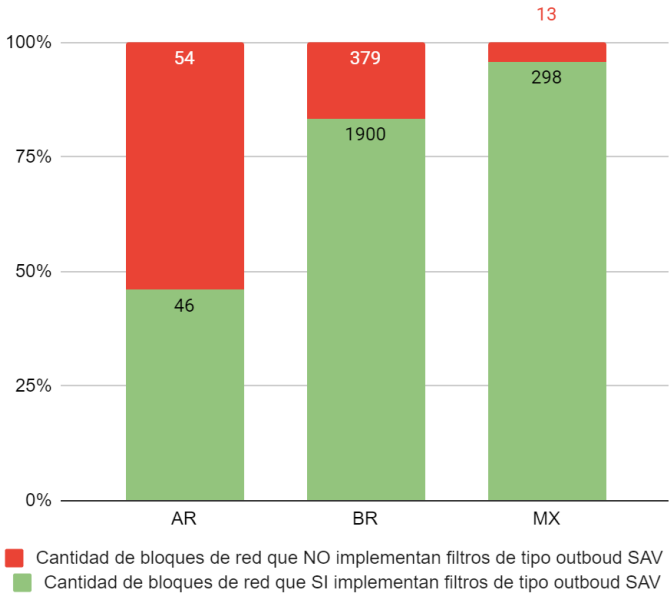

Figure 10 - Outbound IP spoofing in Argentina, Brazil, and Mexico

Figure 10 shows the three countries where the highest number of tests were carried out (more than 100 tests). As the image shows, outbound SAV are implemented in 46% of the blocks tested in Argentina. Likewise, outbound SAV is implemented in 83% of the total number of blocks tested in Brazil and 96% of those tested in Mexico.

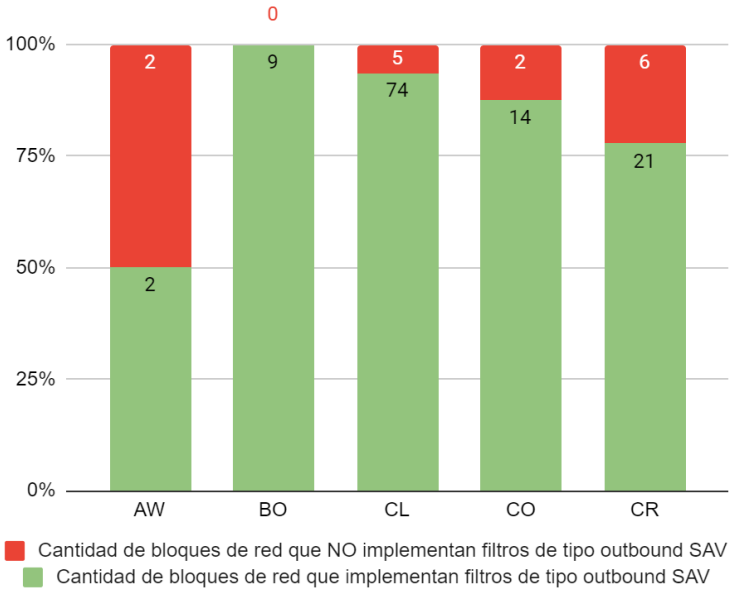

Figure 11 - Outbound IP spoofing in Aruba, Bolivia, Chile, Colombia, and Costa Rica

Figure 11 also shows that a large percentage (more than 75%) of the networks tested in Bolivia, Chile, Colombia, and Costa Rica implement outbound SAV. The percentage of IP blocks that implement SAV filtering is only lower in Aruba (just two or 50% of the tested networks).

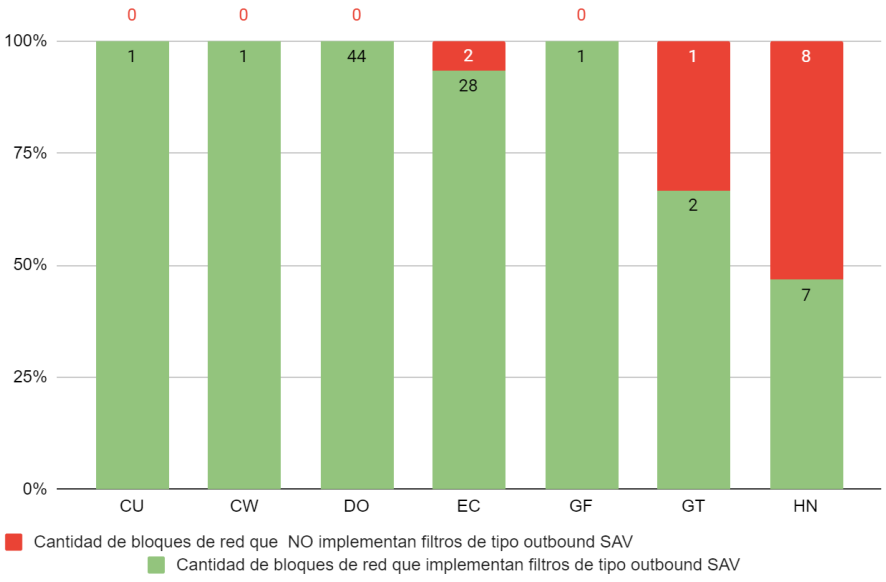

Figure 12 - Outbound IP spoofing in Cuba, Curaçao, Dominica, Ecuador, French Guiana, Guatemala, and Honduras

Figure 12 shows that all IP blocks tested in Cuba, Curaçao, Dominica, and French Guyana implement outbound SAV. In Ecuador, outbound SAV is implemented in 28 IP blocks, which represents more than 90% of tested networks. Likewise, in Guatemala outbound SAV is implemented in two of the three tested IP blocks (66%), while in Honduras the percentage of implementation is lower, almost 50% (seven of the tested blocks).

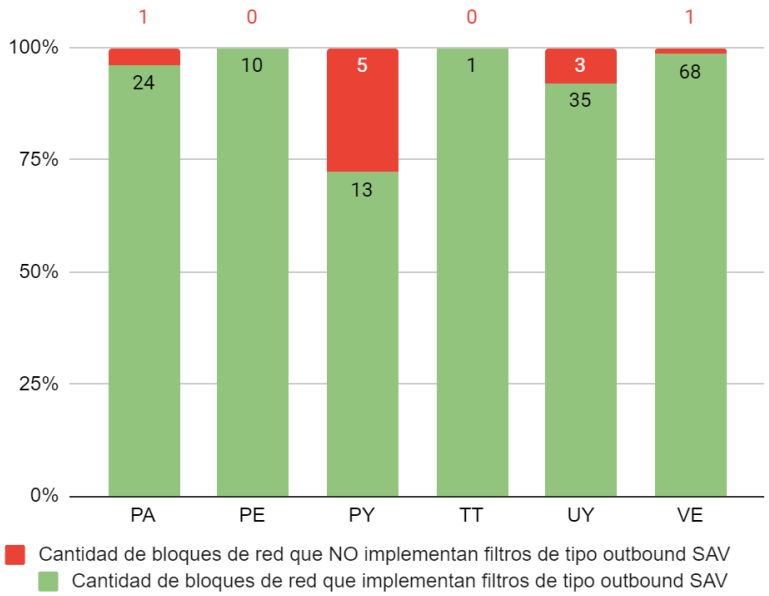

Figure 13 - Outbound IP spoofing Haiti, Panama, Peru, Paraguay, Trinidad and Tobago, Uruguay, and Venezuela

Figure 13 also shows that outbound SAV was implemented in all the networks tested in Peru and Trinidad and Tobago, while in Panama and Venezuela outbound SAV was implemented in almost all tested networks: 24 out of 25 and 68 out of 69, respectively. In Paraguay, 13 out of the 18 tested networks implement outbound SAV. In Uruguay, 35 out of 38. Overall, in these countries, outbound SAV was implemented in more than 70% of the networks, while almost 100% of the blocks tested in four of these countries had deployed SAV.

From a security standpoint, the first step in mitigating a problem is understanding the situation. Based on the information collected in this study, we propose encouraging testing to gain a better understanding of the regional situation in terms of the implementation of inbound and outbound SAV.

In this sense, it is important to define strategies that will allow testing in those countries of the region for which there is no available information and to increase testing in those where only a few networks have been tested.

At the time of writing this report, Brazil is the country where the highest number of tests have been performed, not only in the region but worldwide. A total of 2,279 blocks have been tested in Brazil. This is followed by Mexico with 311 blocks and Argentina with 100 blocks. In the rest of the countries, tests were carried out on less than 100 blocks, with as few as one block tested in some countries.

Inbound SAV deployment

Regarding for the problems related to organizations that allow the reception of spoofed traffic from the Internet, Figure 14 shows the proportion of /24 IPv4 networks per country that were confirmed as being vulnerable by The Closed Resolver Project, as they do not implement inbound SAV.

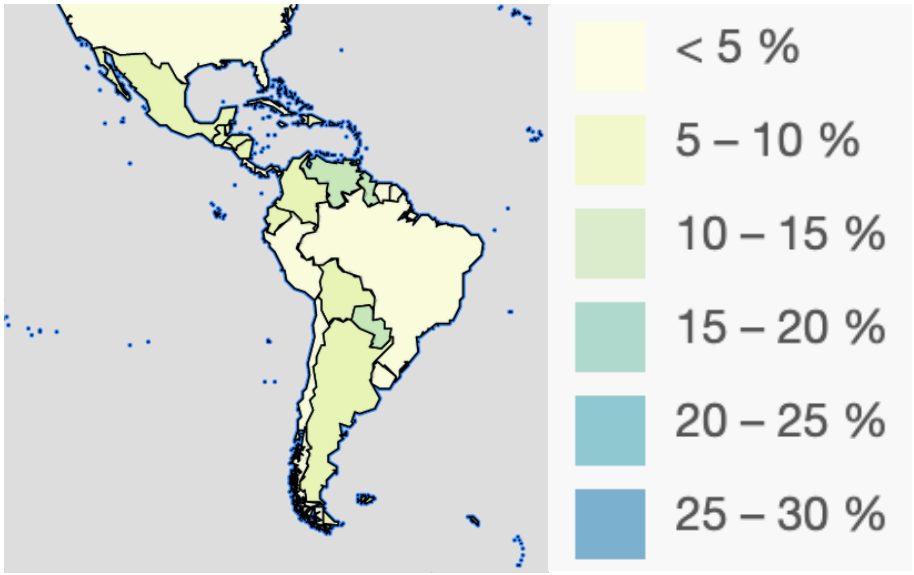

Figure 14 - Percentage of /24 IPv4 networks vulnerable to inbound IP spoofing, by country. ‹https://closedresolver.korlabs.io/#results›, 17 October 2022

According to the information provided by this source, less than 20% of the networks in the region are vulnerable to inbound IP spoofing. Brazil, Chile, French Guyana, Peru, Uruguay, and Suriname are the countries with the lowest percentage of networks vulnerable to inbound IP spoofing, while Guyana, Paraguay, and Venezuela have the highest percentages of networks vulnerable to inbound IP spoofing.

IP spoofing mitigation mechanisms

IP spoofing can be mitigated mainly by applying inbound and outbound SAV. As mentioned earlier, failure to deploy SAV may create different threats. SAV was initially specified in RFC 2827 to address the problem of DDoS attacks that use reflective and amplifying services.

However, it is important to understand that the implementation of this type of filtering also helps to protect the network from other types of attacks, such as those described earlier in this document.

RFC 3704 was created later, considering issues related to the maintenance of the ingress filtering described in RFC 2827 and situations specific to multihomed networks. This RFC proposes different ways to implement SAV. Each will be most appropriate depending on the specific situation. Their effectiveness also varies. Proposed methods include:

- ACL-based SAV.
- Various versions of SAV based on Unicast Reverse Path Forwarding (uRPF) mechanisms. uRPF mechanisms dynamically check whether the source address of a received packet is present in the routing table. Packets whose source address is not present in the routing table are dropped.
    - Strict mode: Checks that each packet's source address is routed through the same network interface that received the packet.
    - Feasible mode: Checks whether any of the routes to the packet's source address (not necessarily the best route) is through the same network interface that received the packet.
    - Loose mode: Verifies the existence of a route to the packet's source address. The interface is immaterial.
    - Loose mode ignoring default routes: Verifies the existence of a route to the packet's source address, but without considering the default routes that may potentially have been configured. The interface is immaterial.

General recommendations

The following recommendations are offered based on those provided by RFC 2827,[25] RFC 3704[26], the MANRS Anti-Spoofing Implementation Guide[27], and the documentation provided by the RIPE Anti-Spoofing Task Force.[28]

The implementation of inbound and outbound SAV policies is recommended to avoid both inbound and outbound IP spoofing. Some of the following techniques may be used:

- ACL-based SAV

  Depending on the number of prefixes, it may be difficult to manually maintain the ACLs implemented by both inbound and outbound SAV.

- SAV filtering based on uRPF Strict mode

  This is an excellent option for networks with symmetric routing where manual ACL maintenance is not required.

- SAV filtering based on uRPF Feasible mode

  According to the MANRS Anti-Spoofing Implementation Guide:

  > Network operators are hesitant to use uRPF in the core of their networks because of the fear of accidentally dropping valid traffic that has taken an unexpected path through their network. uRPF Feasible Path mode should solve such problems.

Configuration tips for Cisco routers

A possible implementation of an ACL in Cisco routers to prevent spoofed traffic is described below.

Figure 15 shows the example included in the MANRS Anti-Spoofing Implementation Guide.[29] In this figure, a red square is used to highlight the ACL that the ISP must implement on the interface to prevent the client from sending spoofed IP traffic to the Internet.

---

[25] ‹https://www.rfc-editor.org/rfc/rfc2827›

[26] ‹https://www.rfc-editor.org/rfc/rfc3704›

[27] ‹https://www.manrs.org/netops/guide/antispoofing/›

[28] ‹https://www.ripe.net/publications/docs/ripe-431›

[29] ‹https://www.manrs.org/netops/guide/antispoofing/

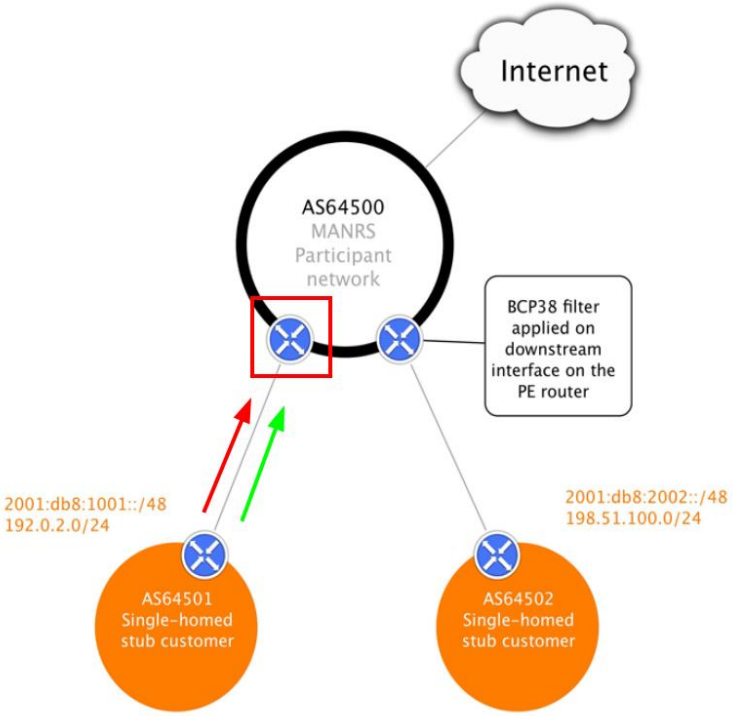

Figure 15 - Sample topology for ACL-based inbound SAV

Assuming that the client's IP blocks are 192.0.2.0/24 and 2001:db8:1001::/48, the ACL that the ISP would configure to avoid receiving spoofed IP traffic would be:

```
    ip access-list extended customer1-in-ipv4
 permit ip 192.0.2.0 0.0.0.255 any
!
ipv6 access-list customer1-in-ipv6
 permit ipv6 2001:db8:1001::/48 any
!
interface x
 ip access-group customer1-in-ipv4 in
     ipv6 traffic-filter customer1-in-ipv6 in
```

To implement SAV based on uRPF Strict mode, filters can be enabled using the following configuration on the network interfaces:

```
    ip cef
    ipv6 cef
    ip verify unicast source reachable-via rx
ipv6 verify unicast source reachable-via rx
```

Configuration tips for Juniper routers

The implementation of an ACL to prevent spoofed traffic in Juniper routers is described below.

Using the example included in the MANRS Anti-Spoofing Implementation Guide,[30] below is the ACL that the ISP can implement to prevent the client from sending spoofed traffic to the internet.

```
    firewall {
    family inet {
    filter customer1-in-ipv4 {
    term allowed-sources {
    from {
    source-address {
    192.0.2.0/24;
    }
    }
    then accept;
    }
    }
    }
    family inet6 {
    filter customer1-in-ipv6 {
```

[30] ‹https://www.manrs.org/netops/guide/antispoofing/›

```
term allowed-sources {
from {
source-address {
2001:db8:1001::/48;
}
}
then accept;
}
}
}
}
interfaces x {
unit 0 {
family inet {
filter {
input customer1-in-ipv4;
}
}
family inet6 {
filter {
input customer1-in-ipv6;
}
}
}
}
```

To implement SAV based on uRPF Strict mode, Juniper router filters can be enabled using the following configuration on the network interfaces:

```
family inet {
```

```
 rpf-check;
}
family inet6 {
 rpf-check;
}
```

To implement SAV based on uRPF Feasible mode in Juniper routers, it must be enabled globally:[31]

```
[edit routing-options forwarding-table]
 unicast-reverse-path feasible-paths;
```

It must also be enabled in the network interface configuration:

```
family inet {
 rpf-check; {
 mode loose;
 }
}
family inet6 {
 rpf-check; {
 mode loose;
 }
}
```

[31] ‹https://www.juniper.net/documentation/us/en/software/junos/interfaces-ethernet-switches/static-routing/topics/ref/statement/unicast-reverse-path-edit-routing-options.html›

Configuration tips for MikroTik routers

In MikroTik routers, configuring an ACL to prevent spoofed traffic from reaching the Internet requires the definition of a list of allowed IP addresses for the organization. This list must consider all the IPs and/or address blocks that may originate packets towards the Internet.

Create a firewall[32] rule that accepts traffic that meets the following conditions:

- *Chain: forward*
- *Src. Address List*: List of allowed IP addresses[33]
- *Out. Interface*: WAN interface that connects to the Internet

Create a firewall rule that drops any traffic that meets the following conditions:

- *Chain: forward*
- Out. Interface: WAN interface that connects to the Internet

To implement SAV based on uRPF Strict mode, MikroTik router filters can be enabled in the IP → Settings → RP Filter = strict menu.[34]

Note: At the time of writing this report, the uRPF Feasible mode is not supported in the latest stable version: RouterOS 7.5.

[32] ‹https://help.mikrotik.com/docs/display/ROS/Filter›

[33] ‹https://help.mikrotik.com/docs/display/ROS/Address-lists

[34] ‹https://help.mikrotik.com/docs/display/ROS/IP+Settings›

Recommendations

The following recommendations are offered to avoid IP address spoofing:

- Assess the status of inbound and outbound SAV in the assigned resources using the tool provided by the CAIDA Spoofer Project.
- Implement inbound and outbound SAV.
- Based on the insights gained from this study, encourage testing and the implementation of best practices recommended for mitigating the issues associated with IP spoofing.
- Use the statistics provided by the CAIDA Spoofer Project to bring visibility to the problem and raise awareness among the community.
- Analyze alternatives to encourage testing and motivate members to join the initiative.

Conclusions

According to the data collected by this study, while most of the tested networks have properly implemented good practices associated with the mitigation of IP spoofing in the region, a part of the universe that was analyzed has not yet acted to mitigate this problem.

This report has unified terminology and organized information from various sources and is intended to be an autonomous and relevant source of information that will allow members to understand the problem of IP spoofing, the mechanisms they can use to assess the status of their networks with respect to IP spoofing, and the actions they can implement to improve their security.

It is expected that this study will lead to a significant increase in the number of networks that apply anti-spoofing techniques, thus contributing to Internet security, stability, and resilience.