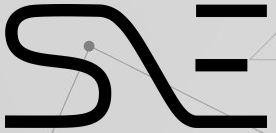


A quantitative analysis of authoritative Domain Name System (DNS) servers and their Resource Public Key Infrastructure (RPKI) adoption

Authors: Sander Post
Brice Habets

Supervisors: Willem Toorop (NLNet Labs)
Tom Carpay (NLNet Labs)

Sunday, 23rd of October, 2022

TABLE OF CONTENTS

INTRODUCTION

Basics

- Domain Name System (DNS) and Border Gateway Protocol (BGP) form the basis of the Internet
- Integrity/security is added on top
  - Resource Public Key Infrastructure (RPKI) for BGP (Integrity)
- RPKI relies on databases, maintained by Regional Internet Registries (RIRs)
- Route Origin Authorisations (ROA) are validated with Route Origin Validation (ROV)

Questions we set out to answer

- “What is the state of RPKI adoption on authoritative name servers?”
  - How many authoritative NS have ROAs and how many return valid?
  - How many authoritative NS validate?
  - How many domains are protected?
  - How many authoritative NS are in both data sets (Duplicates)?

Motivation

- Research into RPKI and ROV adoption of authoritative NS has not been done

BACKGROUND 1/5

Route hijacking

- Maliciously (or unintended) rerouting traffic away from the intended destination
- Multiple ways to hijack[1]
  - We use the traditional sub-prefix attack
    - Announcing a more specific prefix
    - Global impact, should affect the whole network
    - Makes a lot of noise, but no need to be stealthy in this experiment

1) H. Birge-Lee, Y. Sun, A. Edmundson, J. Rexford, and P. Mittal, "Bamboozling Certificate Authorities with BGP," Aug. 2018.

BACKGROUND 2/5

RPKI

- A resource system used to validate received BGP route advertisements
- It is a hierarchical PKI containing prefixes
- The end owner of RPKI is Internet Assigned Numbers Authority (IANA)

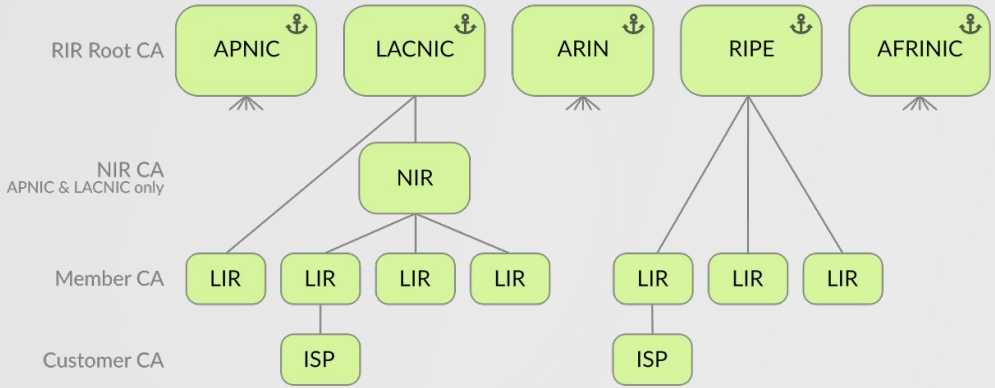

Figure 1: RPKI Hierarchy[2]

2) Rpki.readthedocs.io. 2020. Introduction — RPKI documentation. [online] Available at: https://rpki.readthedocs.io/en/latest/rpki/introduction.html [Accessed 30 June 2022].

BACKGROUND 3/5

RPKI

- Resource certificates are issued
  - The public key in the certificate is bound to the IP address or AS[3]
    - In the certificates IP Address Delegation or AS Identifier Delegation extensions
      - Defined in RFC3779: X.509 Extensions for IP Addresses and AS Identifiers
  - End-Entity (EE) certificates sign resource records, it cannot sign other certificates
  - Owners of a prefix can create ROAs
    - It identifies which prefixes can originate from a given AS

3) M. Lepinski and S. Kent, An Infrastructure to Support Secure Internet Routing, RFC 6480, Feb. 2012. DOI: 10.17487/RFC6480. [Online]. Available: https://www.rfc-editor.org/info/rfc6480.

BACKGROUND 4/5

ROV

- Software called the Relying Party (RP) uses the RPKI infrastructure
  - Four responsibilities[4]
    - Fetching and Caching RPKI Repository Objects
    - Certificate and Certificate Revocation Lists (CRL) Processing
    - Processing RPKI Repository Signed Objects
    - Distributing Validated Cache

4) D. Ma, ZDNS, and S. Kent, Requirements for Resource Public Key Infrastructure (RPKI) Relying Parties, RFC 8897, Sep. 2020. [Online]. Available: https://www.rfceditor.org/rfc/rfc8897.txt.

BACKGROUND 5/5

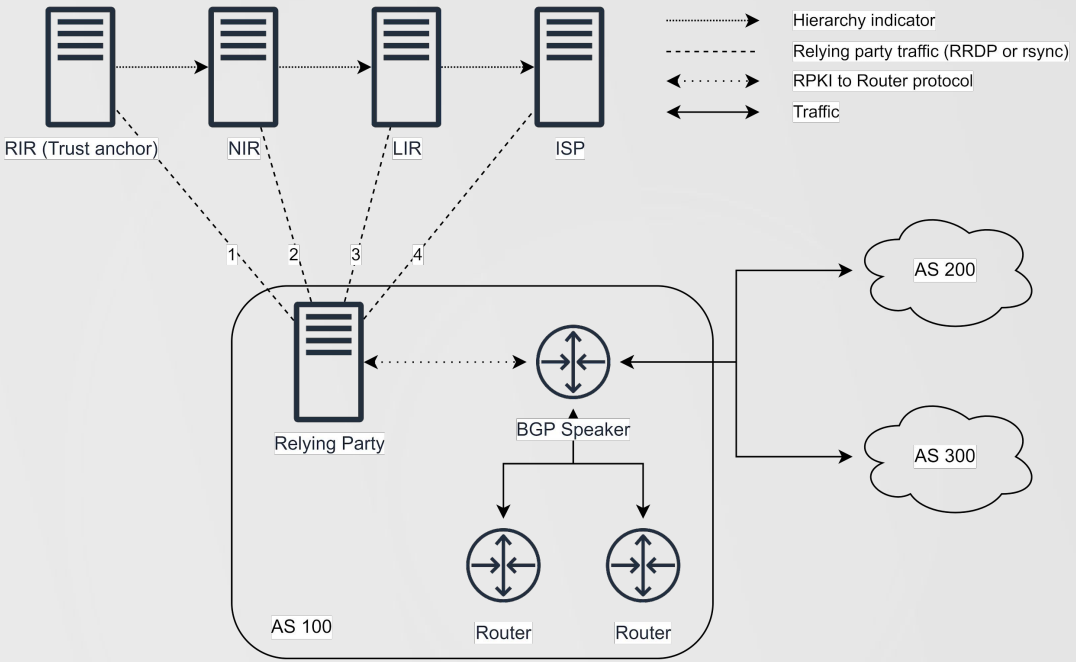

Figure 2: RPKI in practice

RELATED WORK

- In 2020, Linssen tried to measure the protection of 11 top level domains including .com. They concluded that 45% of the domains were covered by a ROA.[5]
  - Research into coverage by ROA's.
  - We also include validation.

5) R. Linssen, "Vulnerability of dns name servers against bgp hijacking," Feb. 2020

METHOD: Environment setup

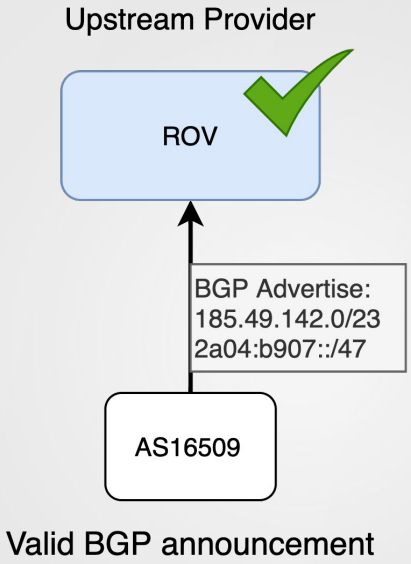

Figure 3: Environment, valid BGP announcements: /23 and /47

METHOD: Environment setup

Figure 4: Environment, added invalid BGP announcement: a /24 and /48

METHOD: Environment setup

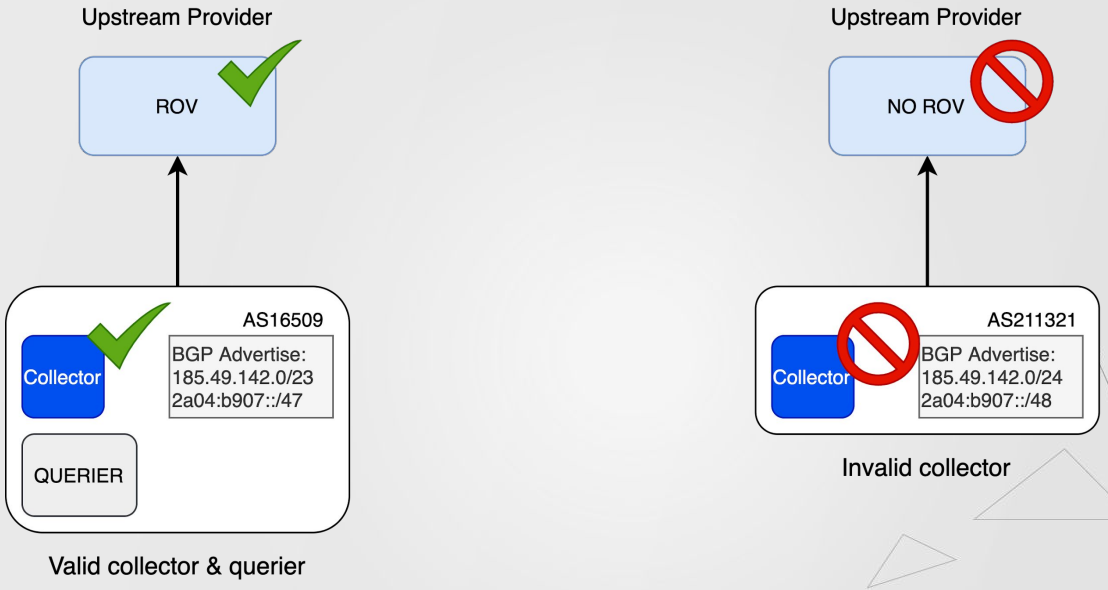

Figure 5: Environment, added querier on valid collector

METHOD: Environment

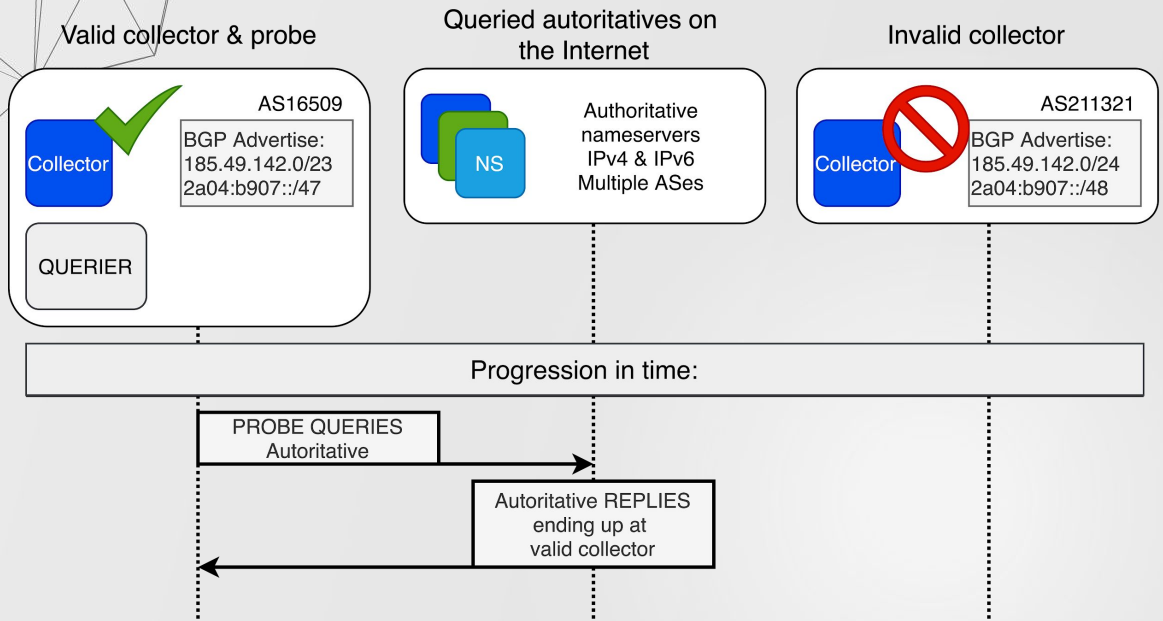

Figure 6: Environment, query response to valid collector

METHOD: Environment

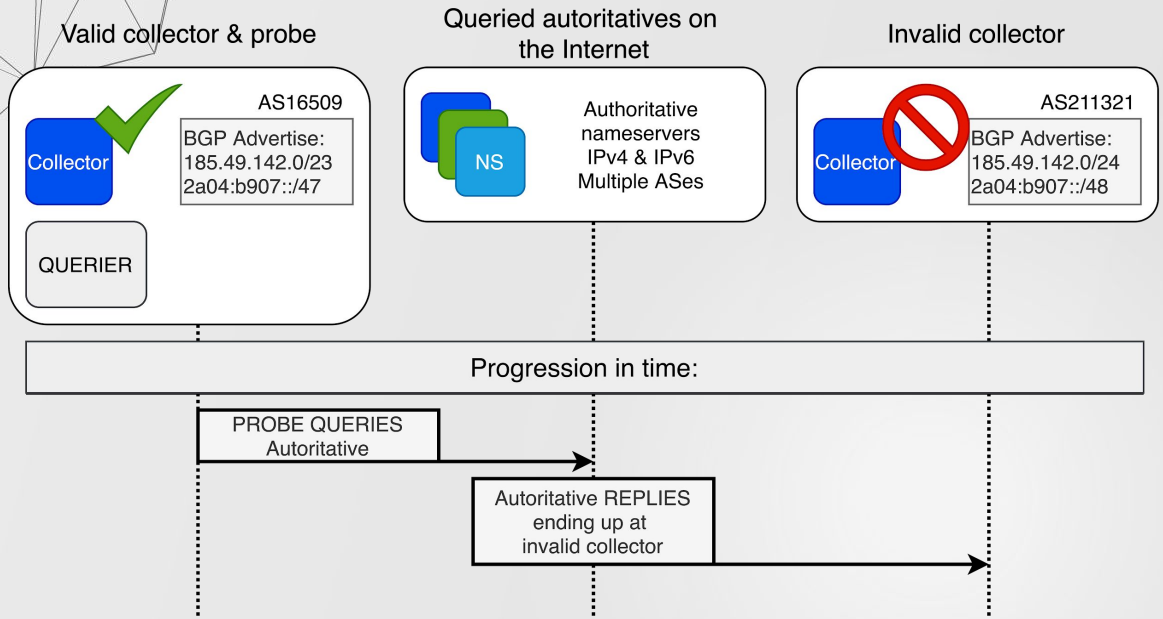

Figure 7: Environment, query response to invalid collector

METHOD: Experiments

Experiments

- Data from OpenINTEL Active DNS Measurements Joint Project[6]
  - List of authoritative NS, both IPv4 and IPv6
  - Generic Top Level Domains (gTLDs)
  - Country-code (ccTLDs)
  - Alexa top 1 million
  - Cisco Umbrella top 1 million
- Experiment 1: sorting the list
- Experiment 2: randomizing the list
- Both lists (IPv4 and IPv6) are queried once per hour

6) Latest news. [Online]. Available: https://openintel.nl/.

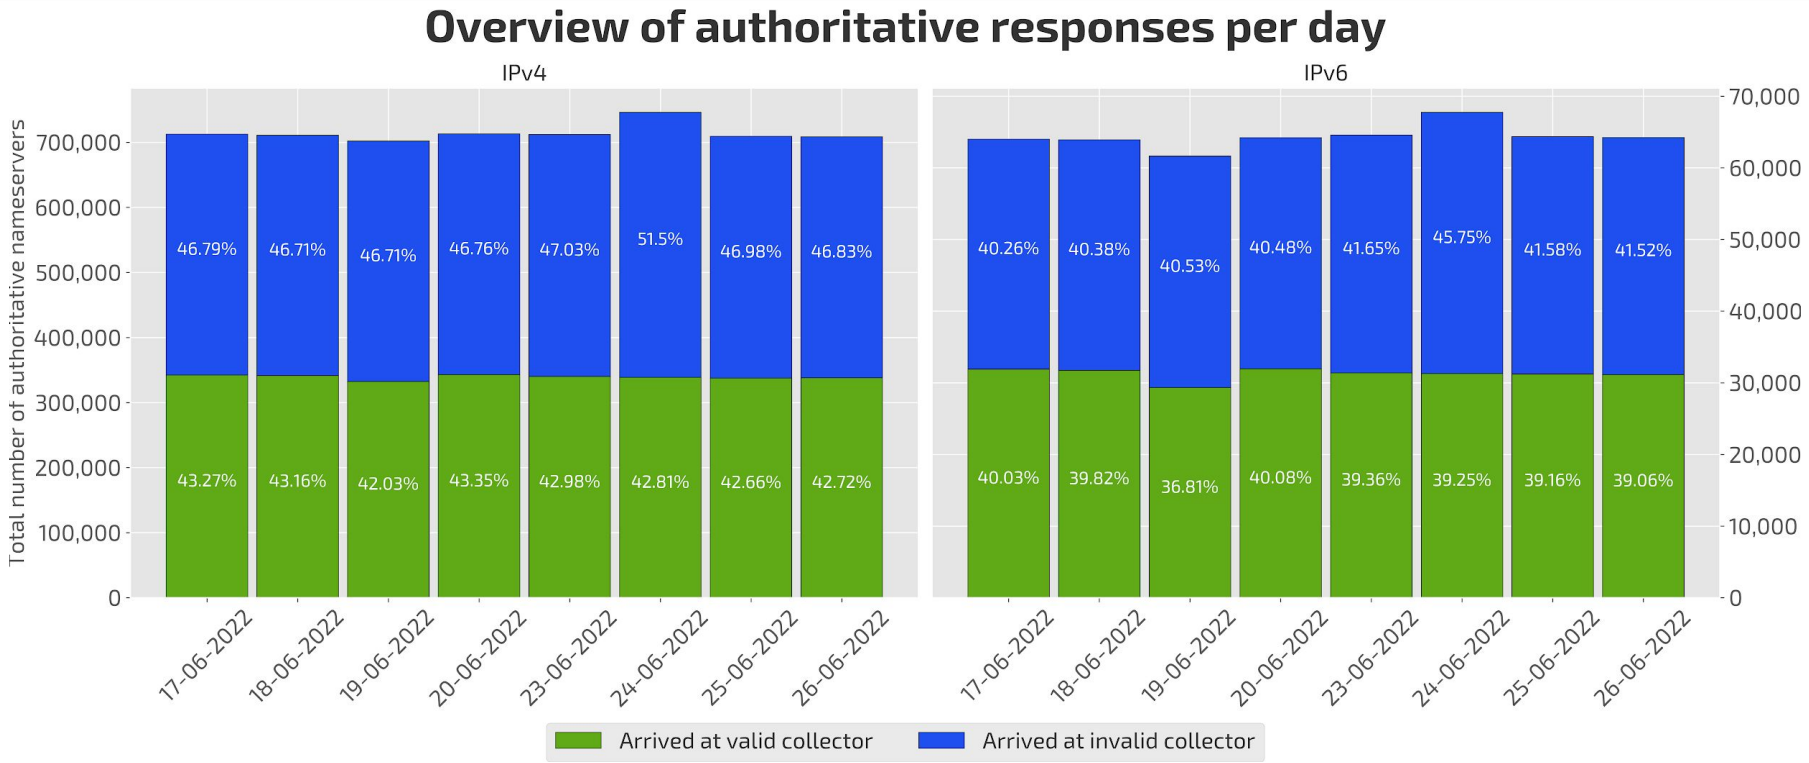

Figure 8: General overview of experiment results per day. On the y-axis, the total amount of authoritative name servers queried can be seen. Left the result of IPv4 and right the results of IPv6 can be observed. Note that the amount of IPv4 addresses is around 900% larger

- Around 43% of the IPv4 responses arrived at a valid collector
- Around 37% of the IPv6 responses arrived at a valid collector

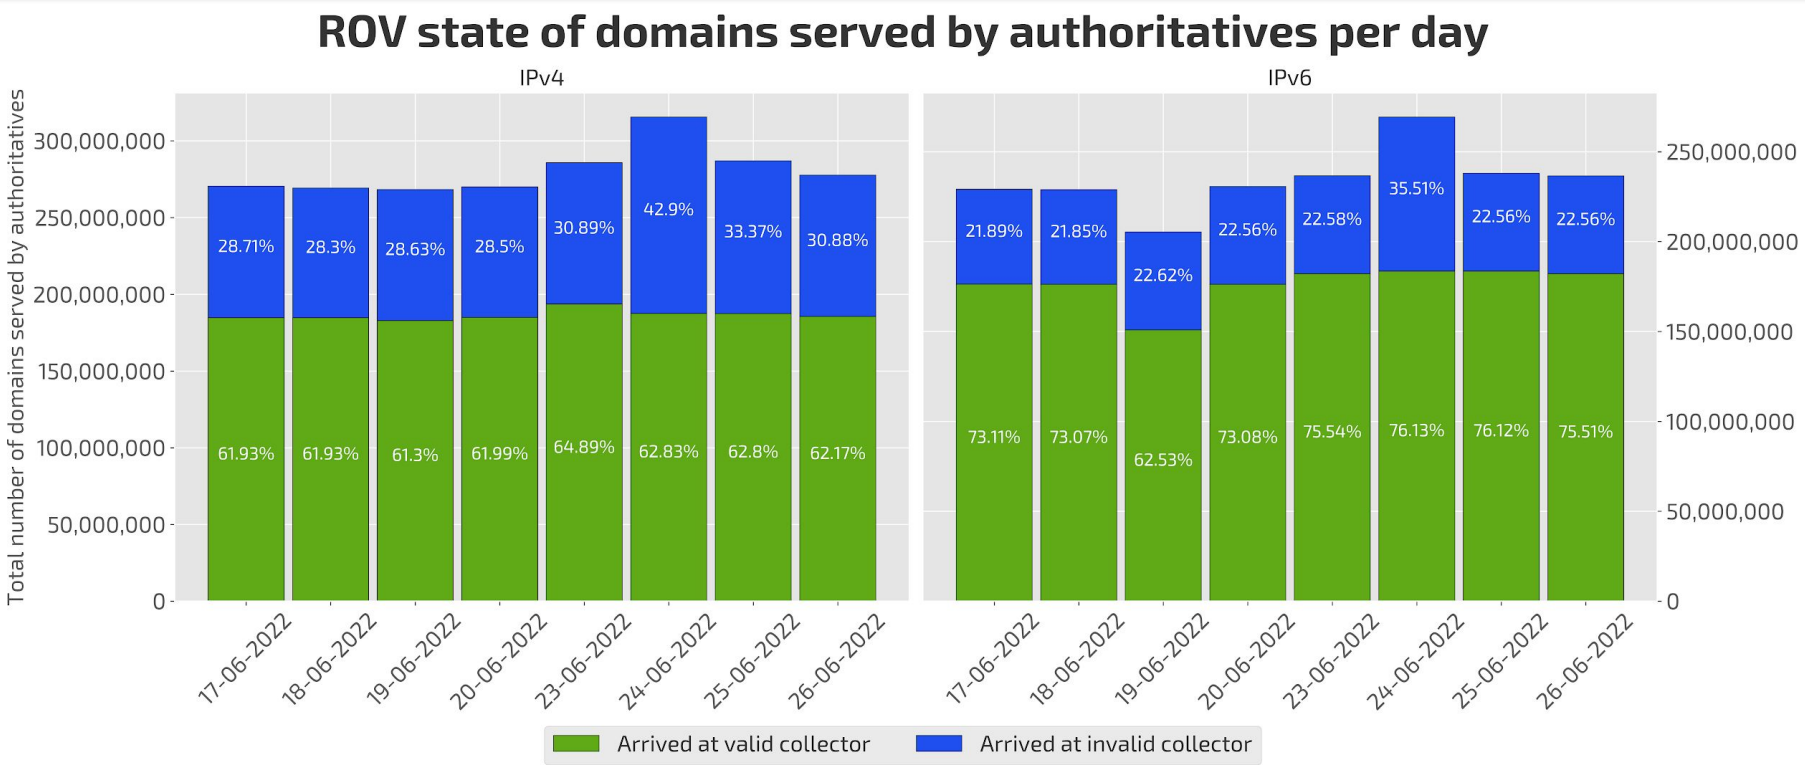

Figure 9: General overview of the amount of domains served by an authoritative name server. On the y-axis we can see the amount of domains. Note that there are way more domains reachable over IPv4 then over IPv6.

- The total number of IPv4 reachable domains is larger
- Proportionally, in essence, more IPv6 reachable domains are protected than IPv4 reachable domains

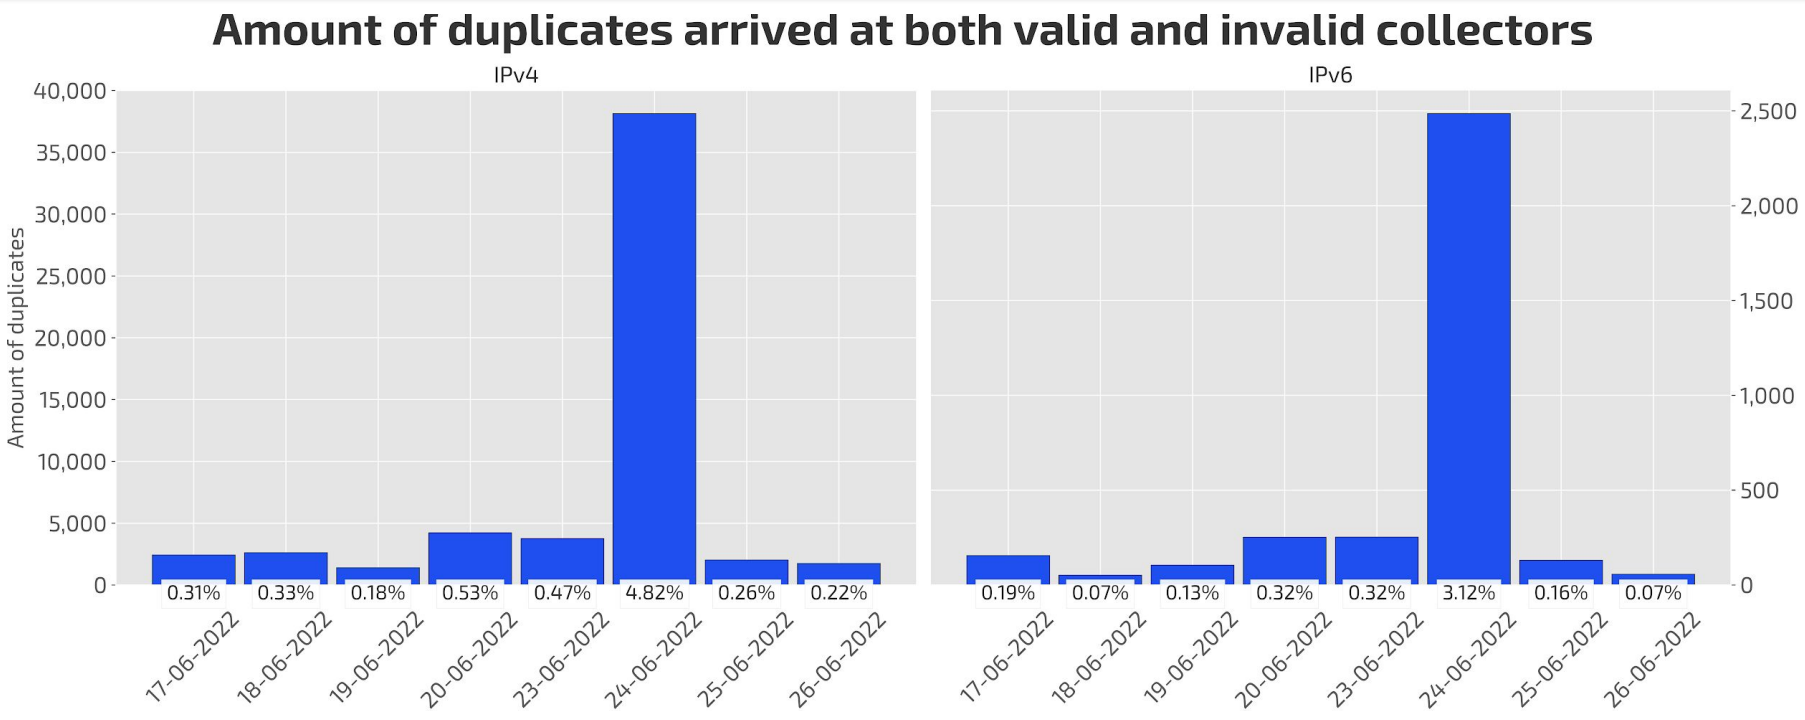

Figure 10: This figure shows the total amount of duplicates. These are responses from authoritative name servers seen on both valid and invalid collectors during different hours. Below the bars, the percentage of the total queried authoritative name servers can be seen.

- On June the 24th there is a huge outlier where responses arrive on both collectors during the day

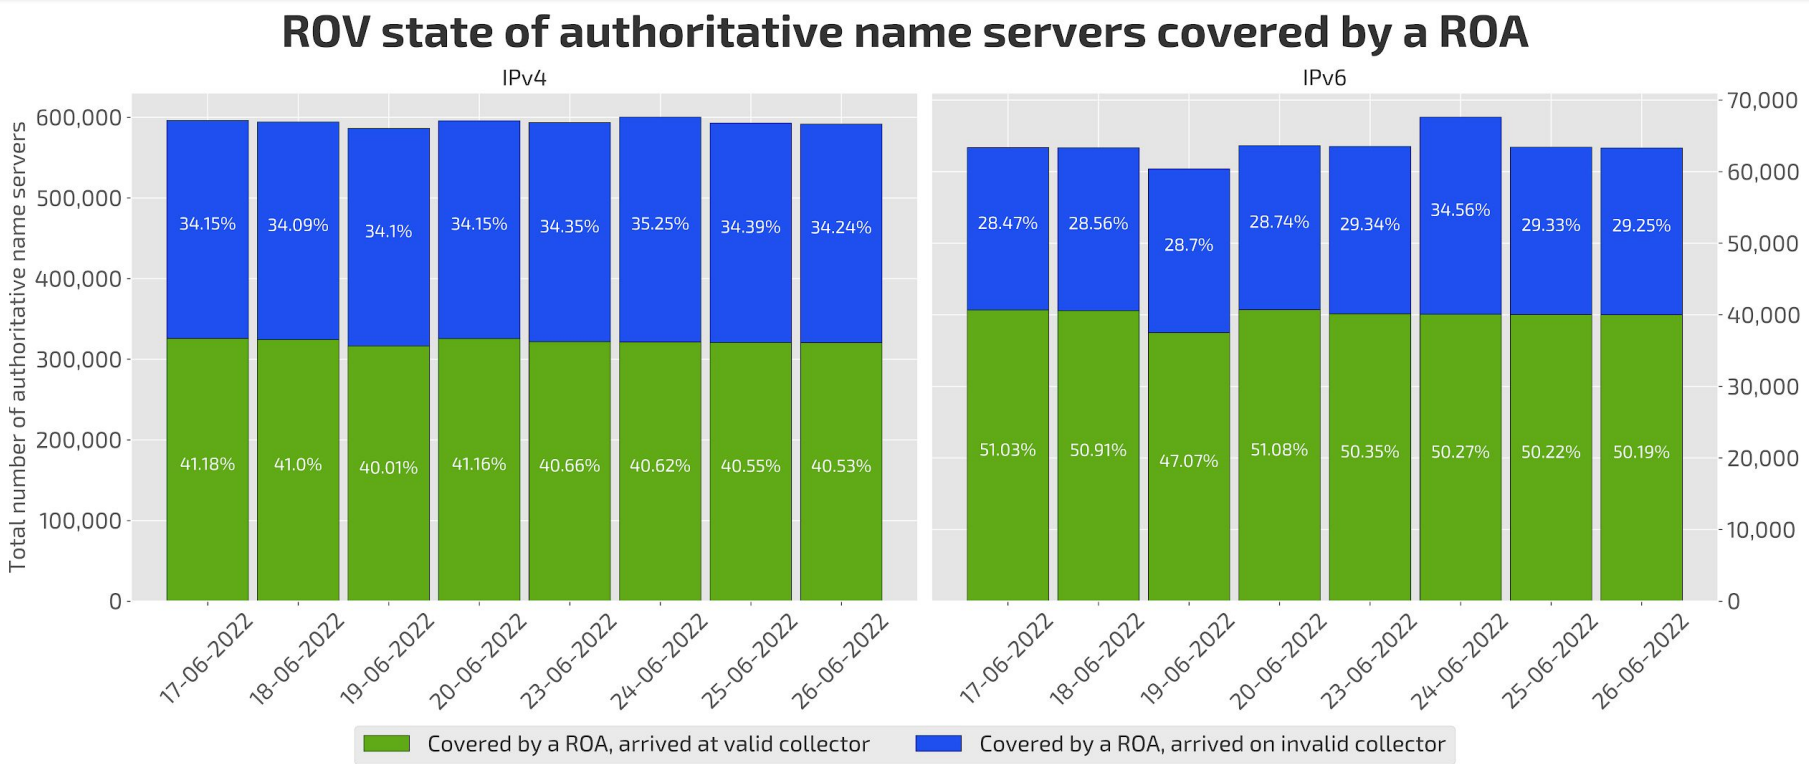

Figure 11: General overview of authoritative name server addresses covered by a ROA. On the y-axis the total amount of addresses is presented. Note that the total amount of IPv4 addresses is once again much larger.

- However, IPv6 addresses are for around 78% covered by ROAs, for IPv4 this is 74%
- Around 40% of the IPv4 responses arrive at the valid collector, for IPv6 this is around 50%
- This could imply that the AS's where the authoritative name servers reside drop invalids (however…)

DISCUSSION: Impact

- We presented an overview of the current state of authoritative name servers and their RPKI adoption
  - Based on a specific dataset
  - If our dataset matches previous research, we see an increase of 66.67% of authoritative name servers protected by a ROA
- How reliable is our data?
  - The Internet is dynamic
  - We have our collectors and probe in one specific place, one point of view
  - Dependent on what happens on the path towards the authoritative, and what happens on the way back?

Discussion: Weakest-link problem

- What if our validating AS is surrounded by non-validating AS's
  - Let's call this the weakest-link problem

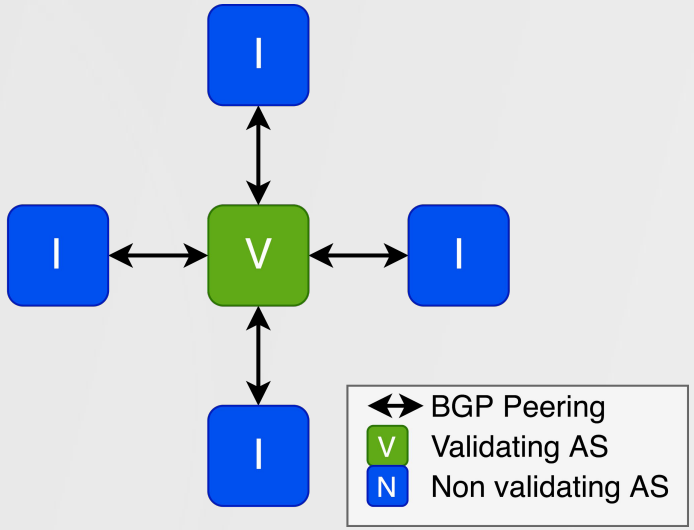

Figure 12: Validating AS surrounded by non validating AS's

Discussion: Weakest-link problem

- What if our validating AS is surrounded by non-validating AS's
  - Let's call this the weakest-link problem
- The upstream decides where the response will end up.
  - Even if the authoritative name server resides in an AS that drops invalid announcements

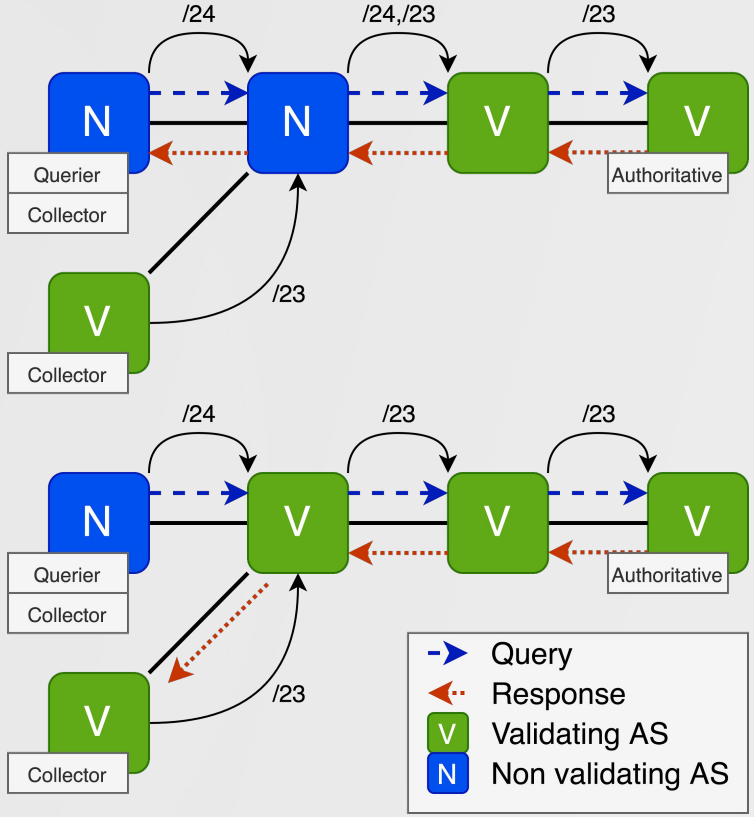

Figure 13: More context to specific AS PATHS

Discussion: Weakest-link problem

- What if our validating AS is surrounded by non-validating AS's
  - Let's call this the weakest-link problem
- The upstream decides where the response will end up.
  - Even if the authoritative name server resides in an AS that drops invalid announcements
- Now we have a valid upstream, but the authoritative name server AS doesn't drop invalid announcements
  - The response arrives at the valid collector
  - The upstream decides where the response arrives
  - Even intermediate non validating hops will contribute to this outcome

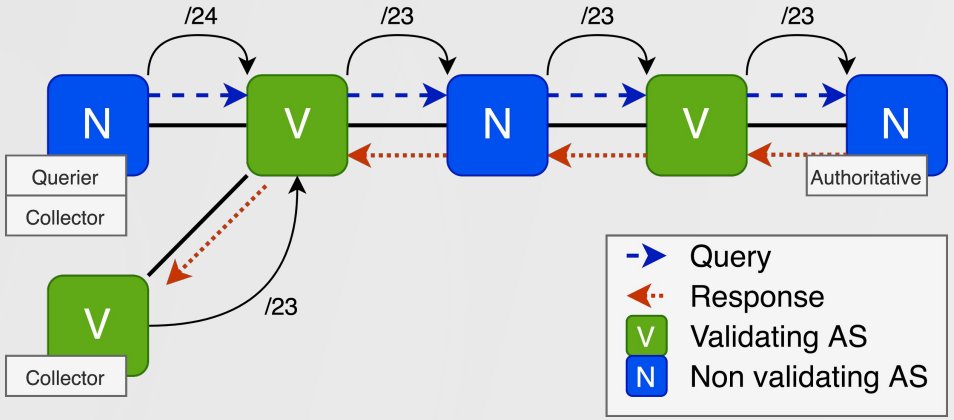

Figure 14: Does it really matter if the operator does ROV?

Discussion: Future Work

- Paths on the Internet seem dynamic
  - Research into how dynamic the paths on the Internet really are
  - Do they change often from one point of view or is this an illusion
  - Measuring and comparing paths from the same source to the same destination of an x amount of time
- What is the current state of RPKI adoption
  - Previous research done in 2020, measuring the state of RPKI
  - Introduce the “TraceROV” tool
  - Based on the same mechanics as traceroute and our research, we could measure if operators implement ROV and if they drop invalid announcements

CONCLUSION

- From our experiments we measured that:
  - On average 45% of the queried authoritative name servers reside in an Autonomous System (AS) that is doing Route Origin Validation (ROV)
  - On average 67% of the domains served are measured to reside in an AS doing ROV
- We showed that the Internet is a dynamic place
  - We showed the weakest-link problem
  - We have seen responses on both the valid and invalid collectors
- We can observe that Resource Public Key Infrastructure (RPKI) is not adopted everywhere and possibly still allows for sub-prefix hijacks
  - But adoption is measurably increasing

CONCLUSION

- From our experiments we measured that:
  - On average 45% of the queried authoritative name servers reside in an Autonomous System (AS) that is doing Route Origin Validation (ROV)
  - On average 67% of the domains served are measured to reside in an AS doing ROV
- We showed that the Internet is a dynamic place
  - We showed the weakest-link problem
  - We have seen responses on both the valid and invalid collectors
- We can observe that Resource Public Key Infrastructure (RPKI) is not adopted everywhere and possibly still allows for sub-prefix hijacks
  - But adoption is measurably increasing

RPKI adoption is steadily increasing

Thank you for attending our presentation

NLNETLABS

SIDN

routz

ventus IT PROFESSIONALS

List of OpenINTEL coverage[7]

- ".com"
- ".net"
- ".org"
- ".info"
- ".mobi"
- Around 1200 new gTLDs
  - via the Internet Corporation for Assigned Names and Numbers (ICANN) Centralized Zone Data Service (CZDS)
- ".gov" and ".fed.us"
  - Obtained via the US Federal Government open access API
- ".name"
- ".biz"
- ".asia"
- ".aero"

- ".nl" (The Netherlands)
- ".se" (Sweden)
- ".nu" (Niue)
- ".ca" (Canada)
- ".fi" (Finland)
- ".at" (Austria)
- ".dk" (Denmark)
- ".ru" (Russian Federation)
- ".us" (United States of America)
- ".gt" (Guatemala)
- ".na" (Namibia)
- ".ee" (Estonia)
- ".co" (Colombia)
- ".ch" (Switzerland)
- ".li" (Liechtenstein)
- ".sk" (Slovakia)
- Alexa top 1 million
- Cisco Umbrella top 1 million

7) Current coverage. [Online]. Available: https://openintel.nl/coverage.