The Current State of DNS Resolvers and RPKI Protection

Hero Clipart
© FreePNGImg.com

University of Amsterdam
Erik Dekker
Marius Brouwer

NLnetLabs
Willem Toorop

DNSSEC Workshop at the

ICANN 69 VIRTUAL ANNUAL GENERAL

21 October 2020

Motivation

- DNSSEC protects against address forgery
- But the address can be trivially hijacked

RPKI 101

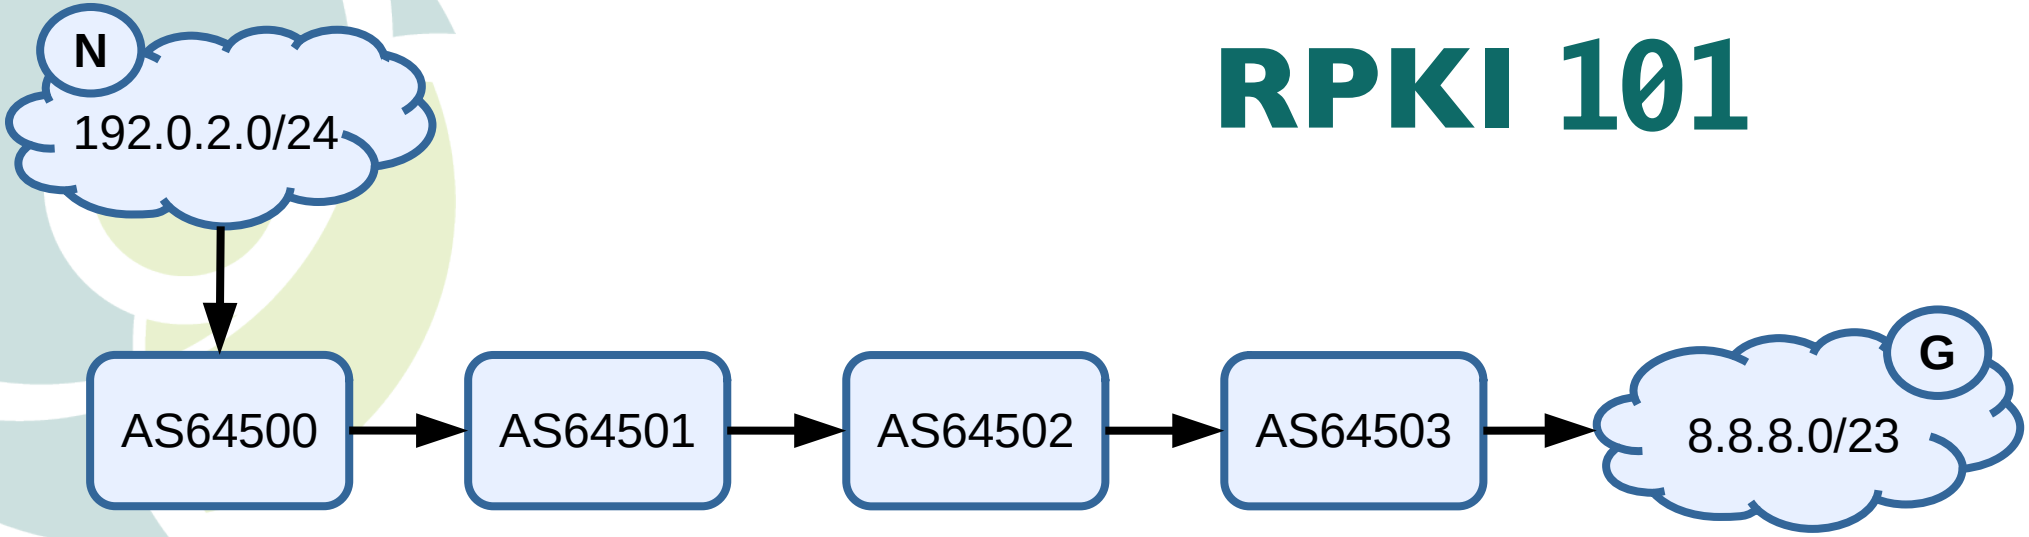

RPKI 101

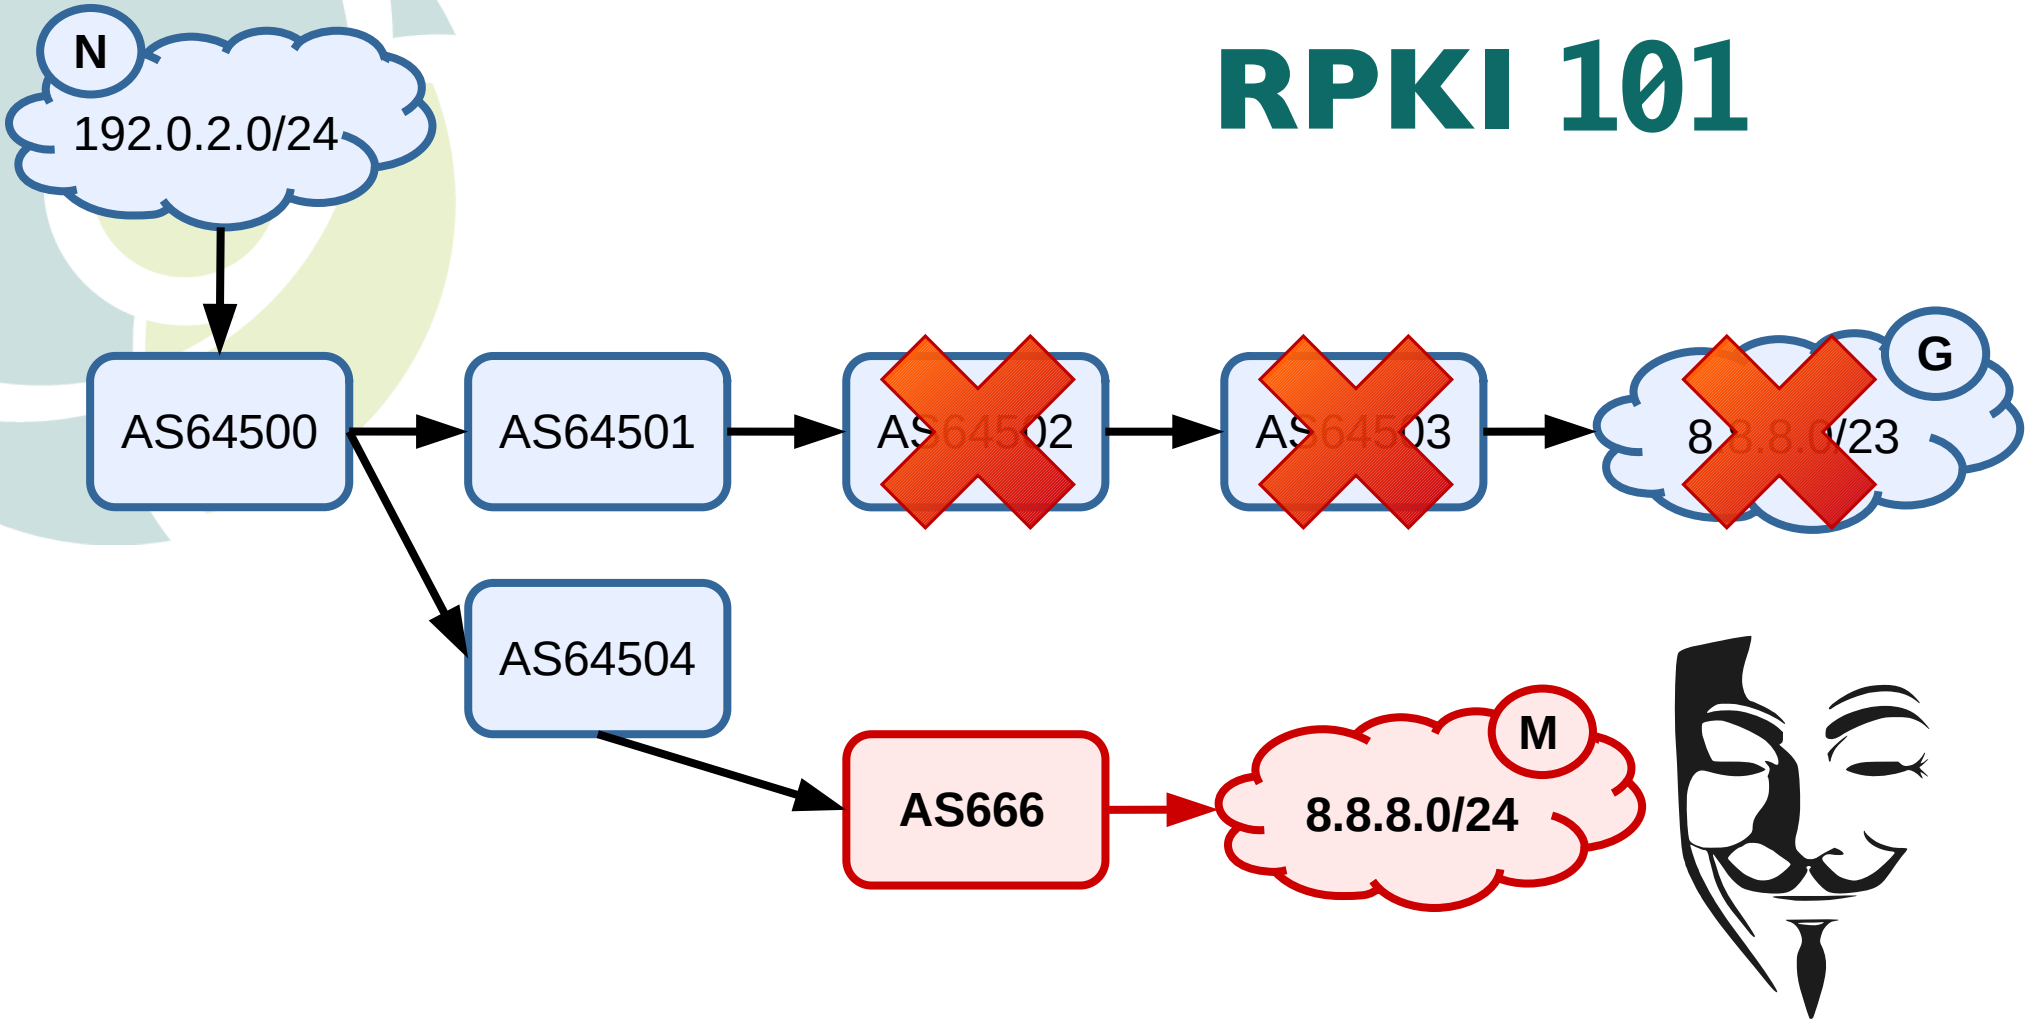

Motivation

- DNSSEC protects against address forgery
- But the address can be trivially hijacked
- RPKI to the rescue

RPKI 101

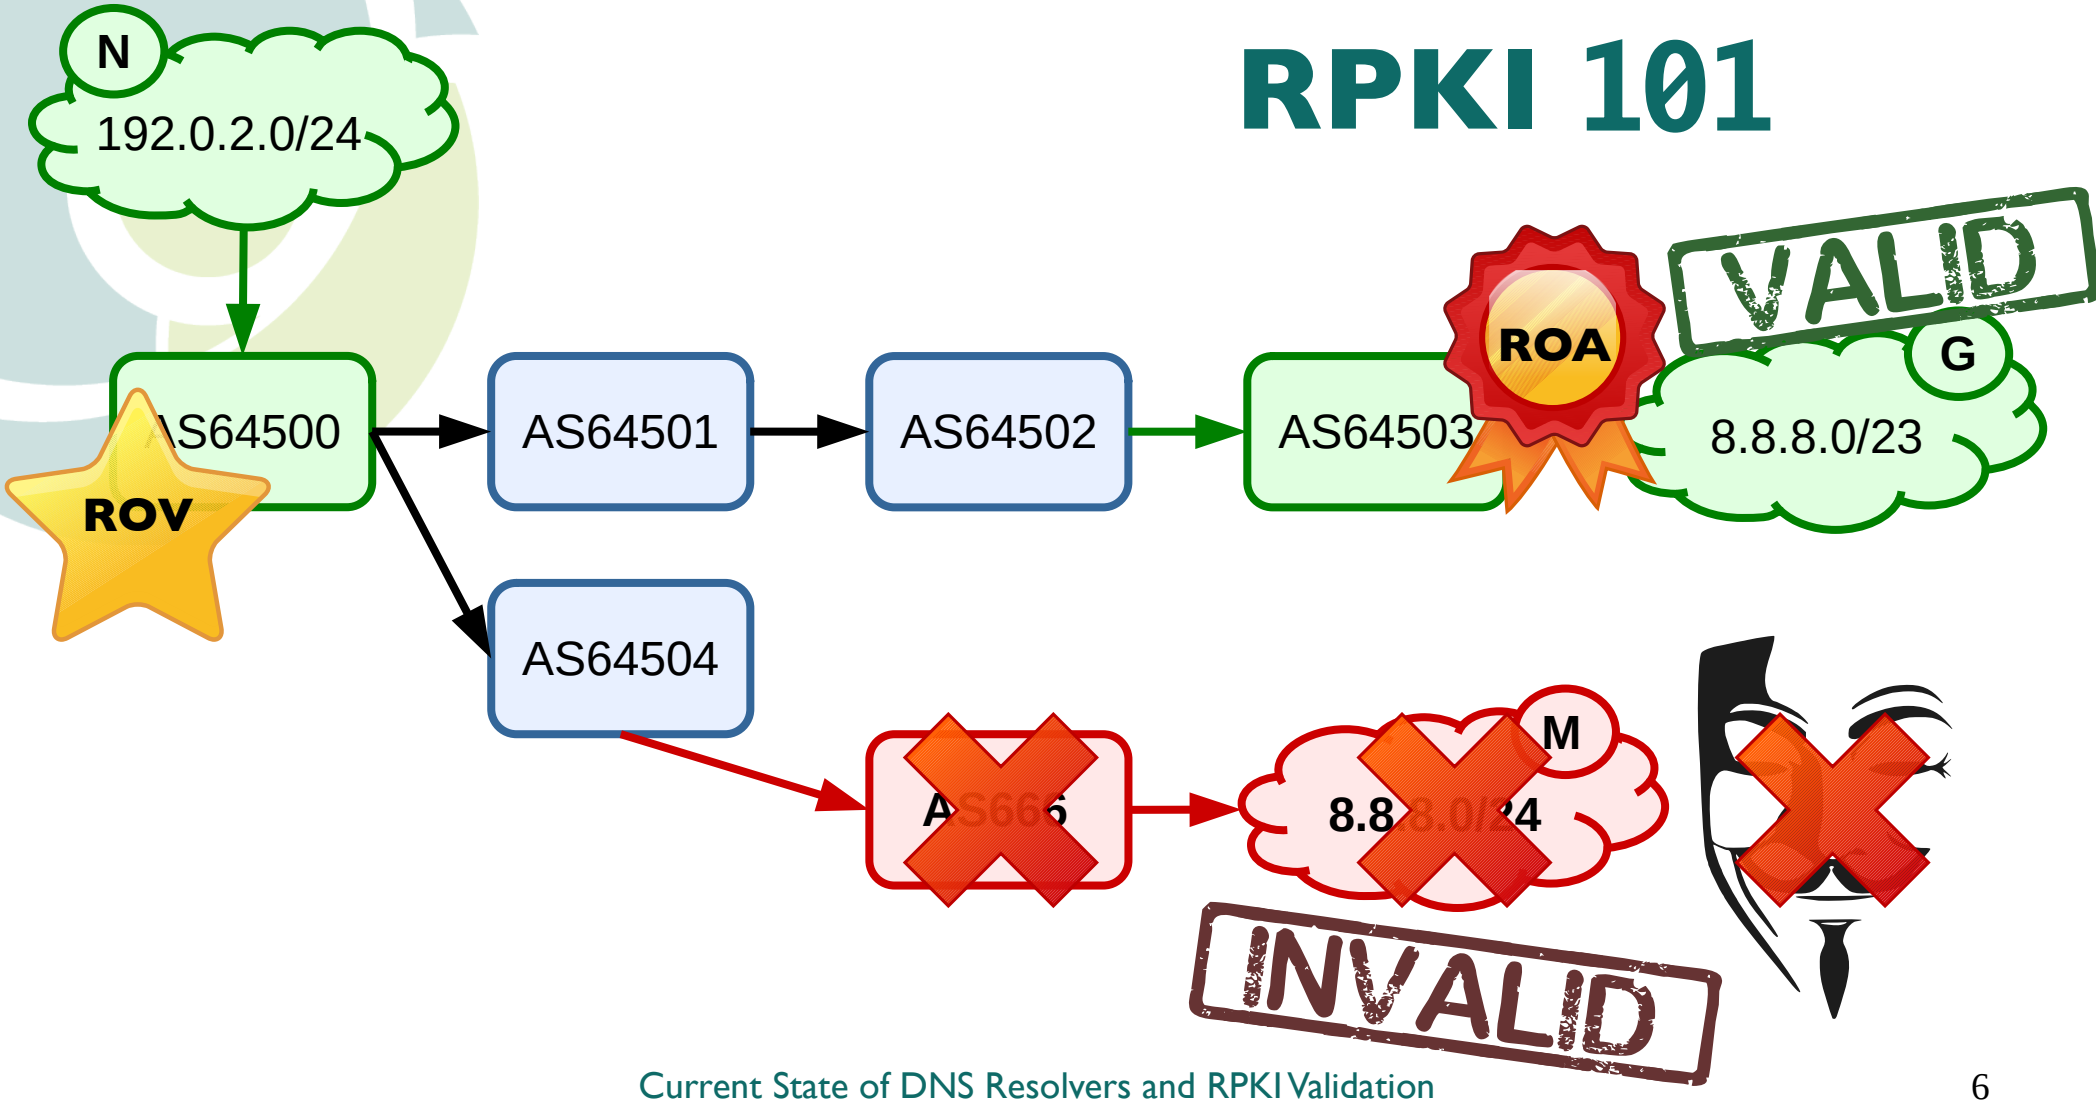
N
192.0.2.0/24
AS64500
ROV
AS64501
AS64502
AS64503
ROA
VALID
G
8.8.8.0/23
AS64504
M
INVALID

Motivation

- What does this have to do with DNS Resolvers?

RPKI 101

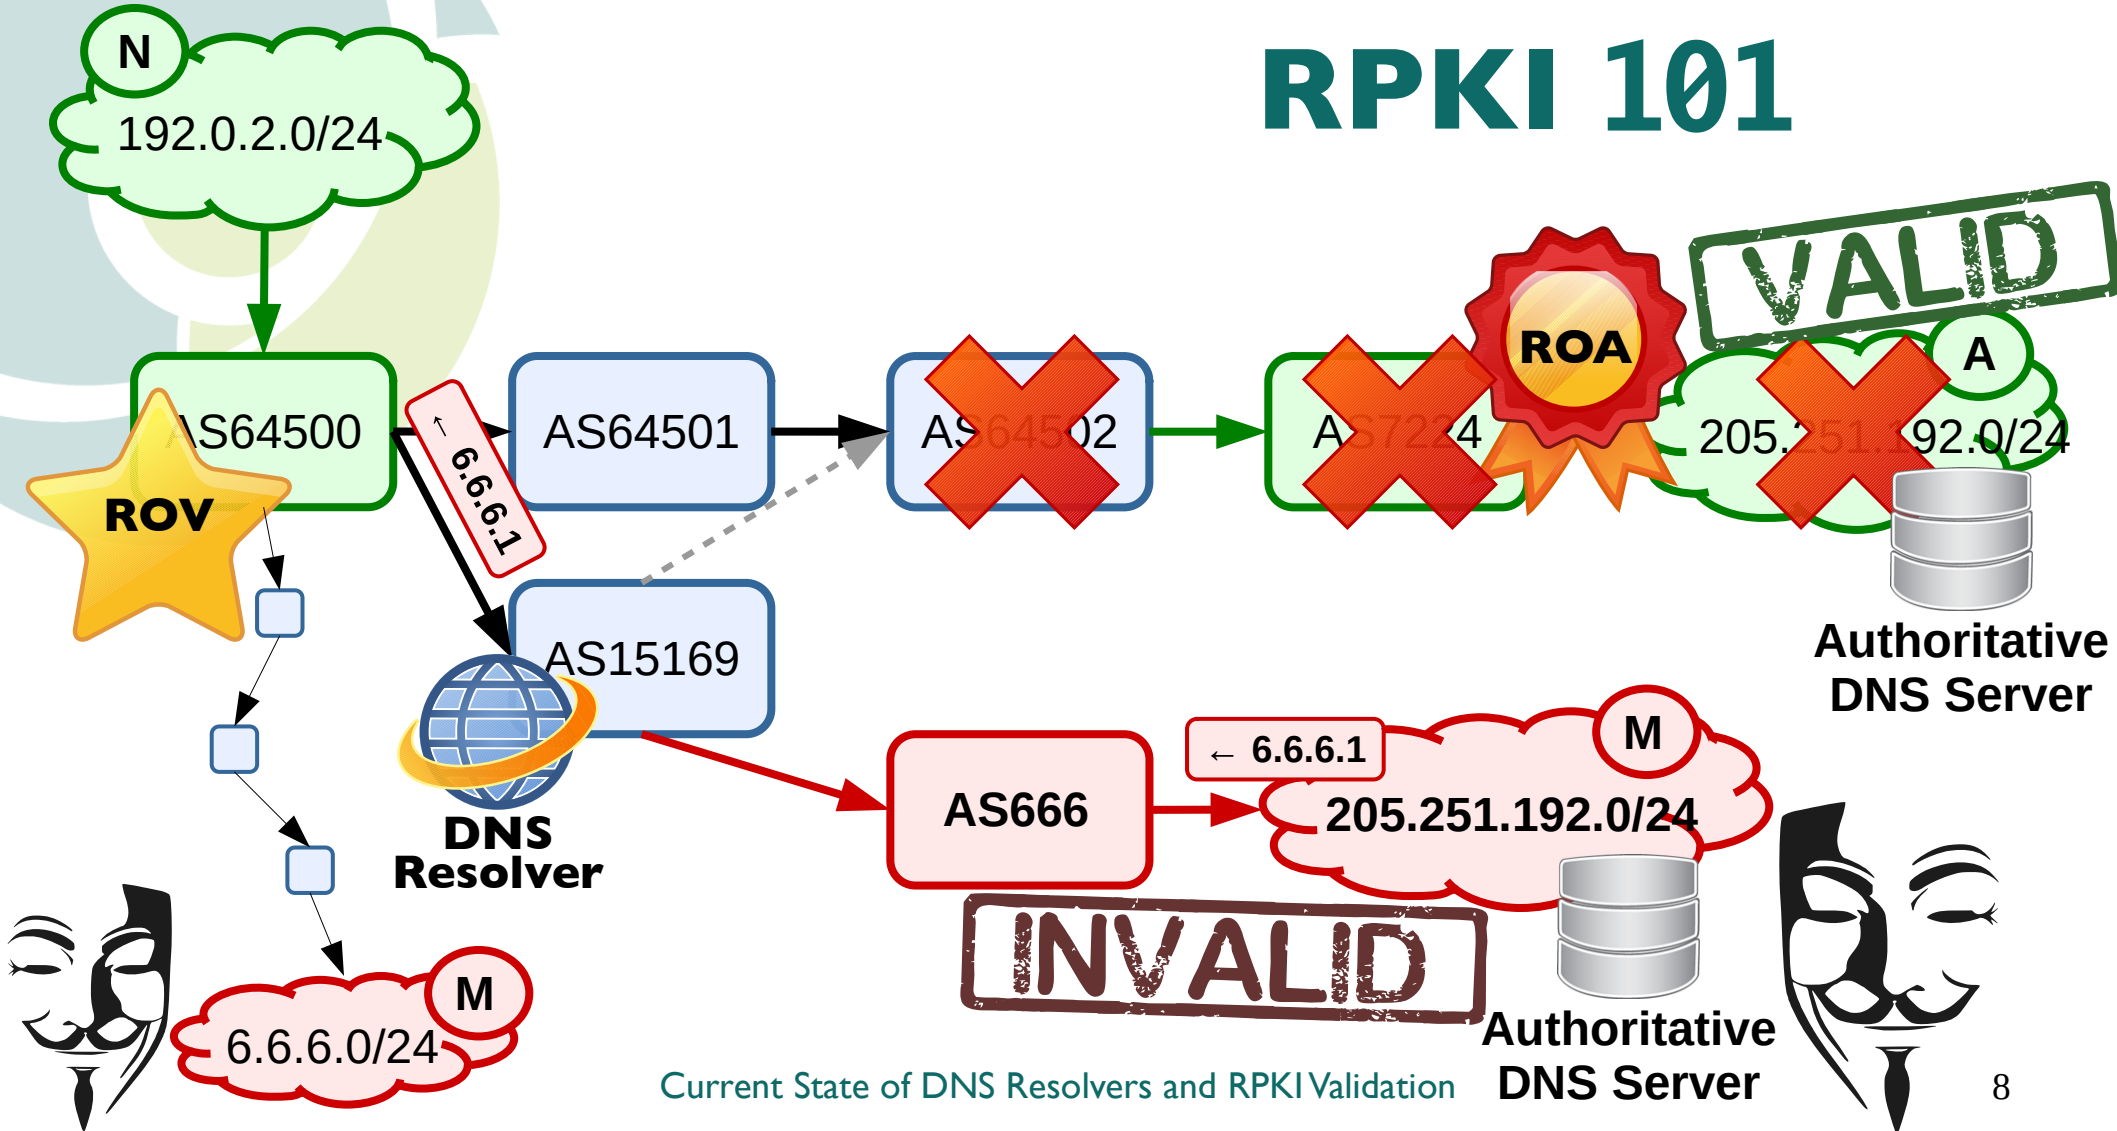
N
192.0.2.0/24
AS64500
ROV
← 6.6.6.1
AS64501
AS64502
AS7224
ROA
VALID
A
205.251.192.0/24
Authoritative DNS Server
AS15169
DNS Resolver
AS666
← 6.6.6.1
M
205.251.192.0/24
INVALID
Authoritative DNS Server
M
6.6.6.0/24


N
192.0.2.0/2
AS64500
ROV
6.6.6.
What Happened? The Amazon Route 53 BGP Hijack to Take Over Ethereum Cryptocurrency Wallets | Internet Society - Chromium
What Happened? The Am
internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/
Internet Society
Mutually Agreed Norms for Routing Security (MANRS) 27 April 2018
EN ES
What Happened? The Amazon Route 53 BGP Hijack to Take Over Ethereum Cryptocurrency Wallets
By Aftab Siddiqui
Senior Manager, Internet Technology - Asia-Pacific
Yesterday, we published a blog post sharing the news and some initial details about Amazon's DNS route hijack event to steal Ethereum cryptocurrency from myetherwallet.com. In this post, we'll explore more details about the incident from the BGP hijack's perspective.
ALID
A
92.0/24
Authoritative
DNS Server

RPKI vs DNSSEC

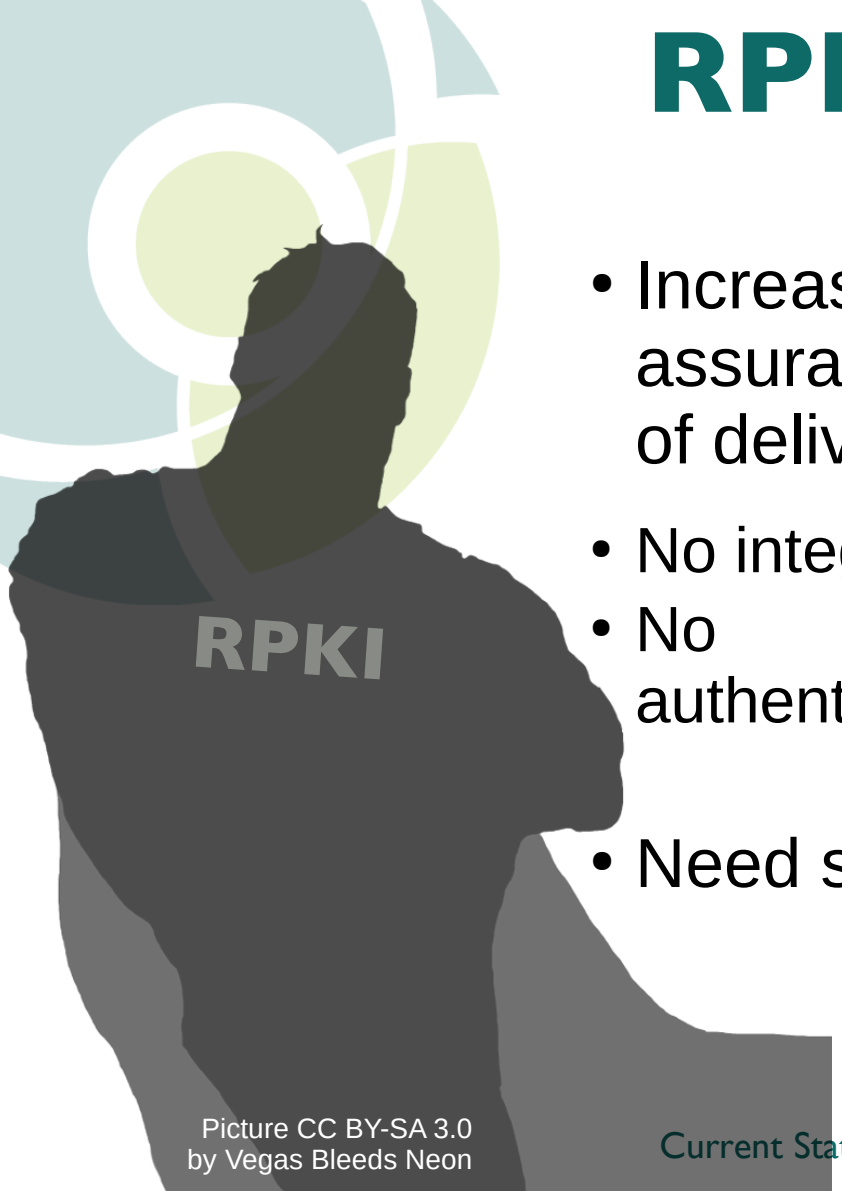

- Increase assurance of delivery
- No integrity
- No authentication

- Doesn't matter how you got it
- Integrity
- Origin authentication

- Need signing **and** validation

Research question

Main:

What is the state of Route Origin Validation (RoV) on DNS resolvers?

Sub:

- Does the length of the AS path matter?
- How does anycast influence the protection?

Test setup

```
$ORIGIN rootcanary.net
$TTL 60
@     SOA  ns1.surfnet.nl. (
           dns-beheer.surfnet.nl.
           2020080503 ; serial
           10800      ; refresh
           3600       ; retry
           604800     ; expire
           86400      ; minimum
           )
      NS   ns1.surfnet.nl.
      NS   ns2.surfnet.nl.
      NS   ns3.surfnet.nl.
      NS   ns1.zurich.surf.net.

$TTL 25200

valid4     NS   valid4
valid4     A    209.24.1.6

invalid4   NS   invalid4
invalid4   A    194.32.71.6
```

```
$ORIGIN valid4.rootcanary.net
$TTL 300
@          SOA  valid4.rootcanary.net. (
                sysadm.rootcanary.org.
                2020012100 10800 3600
                604800 300 )
           NS   @
           A    209.24.1.6
$TTL 1
invalid    DNAME invalid4.rootcanary.net.
```

prefix	209.24.1.0/24
max len	24
ASN	15562


```
$ORIGIN invalid4.rootcanary.net
$TTL 300
@          SOA  invalid4.rootcanary.net. (
                sysadm.rootcanary.org.
                2020012100 10800 3600
                604800 300 )
           NS   @
           A    194.32.71.6
*          A    145.97.20.20
```

prefix	194.32.71.0/24
max len	24
ASN	0

Test setup

RIPE NCC
RIPE Atlas

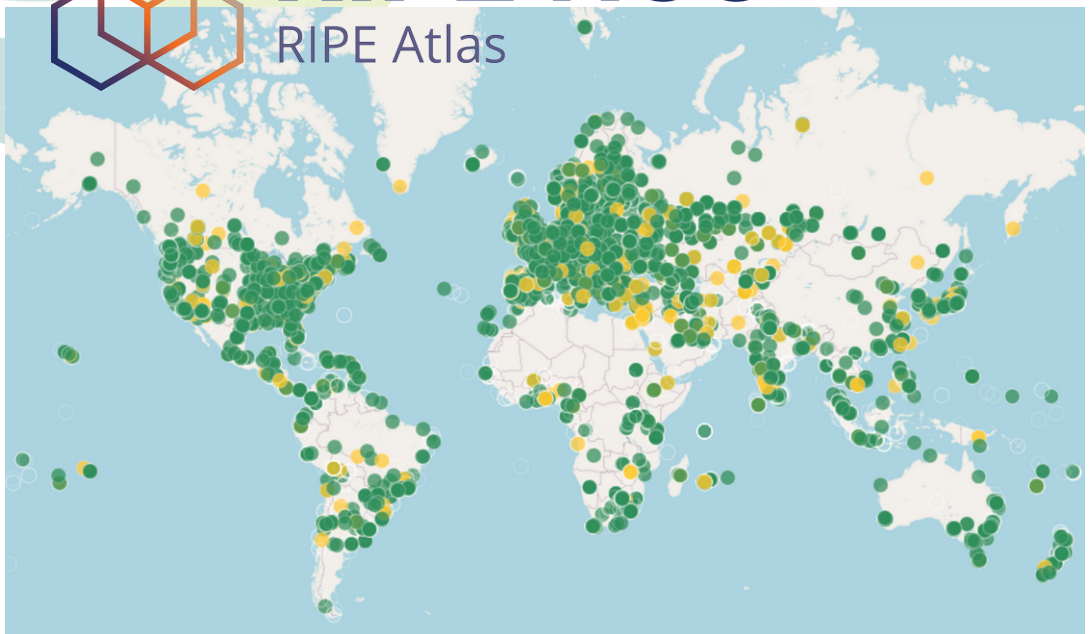

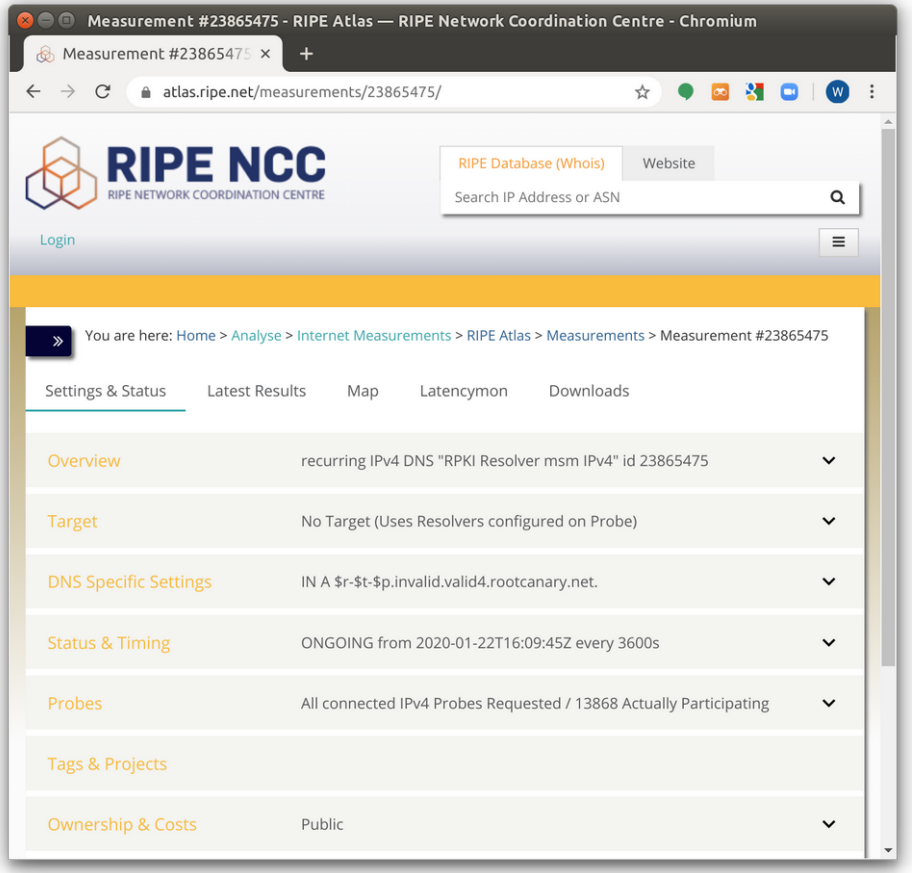

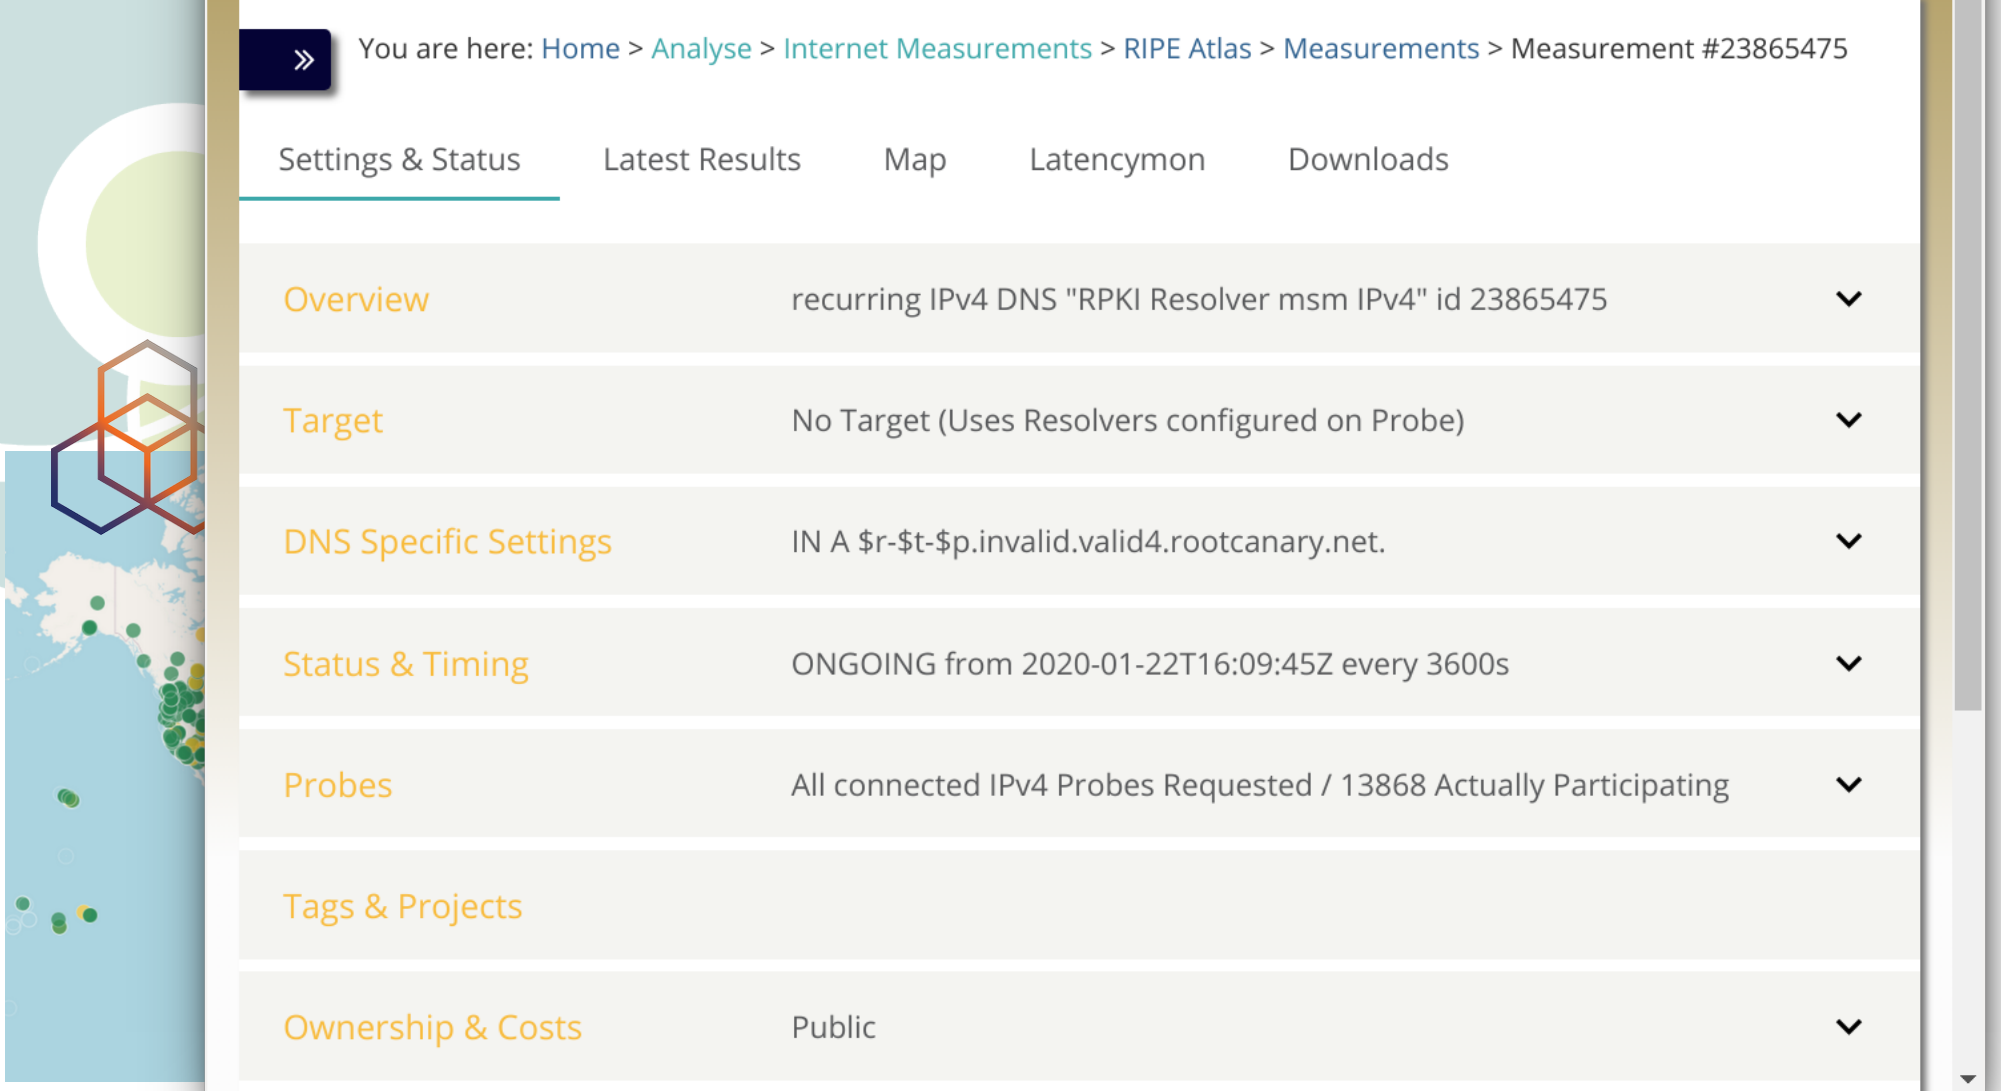
You are here: Home > Analyse > Internet Measurements > RIPE Atlas > Measurements > Measurement #23865475
Settings & Status
Latest Results
Map
Latencymon
Downloads
Overview
recurring IPv4 DNS "RPKI Resolver msm IPv4" id 23865475
Target
No Target (Uses Resolvers configured on Probe)
DNS Specific Settings
IN A $r-$t-$p.invalid.valid4.rootcanary.net.
Status & Timing
ONGOING from 2020-01-22T16:09:45Z every 3600s
Probes
All connected IPv4 Probes Requested / 13868 Actually Participating
Tags & Projects
Ownership & Costs
Public

Test setup

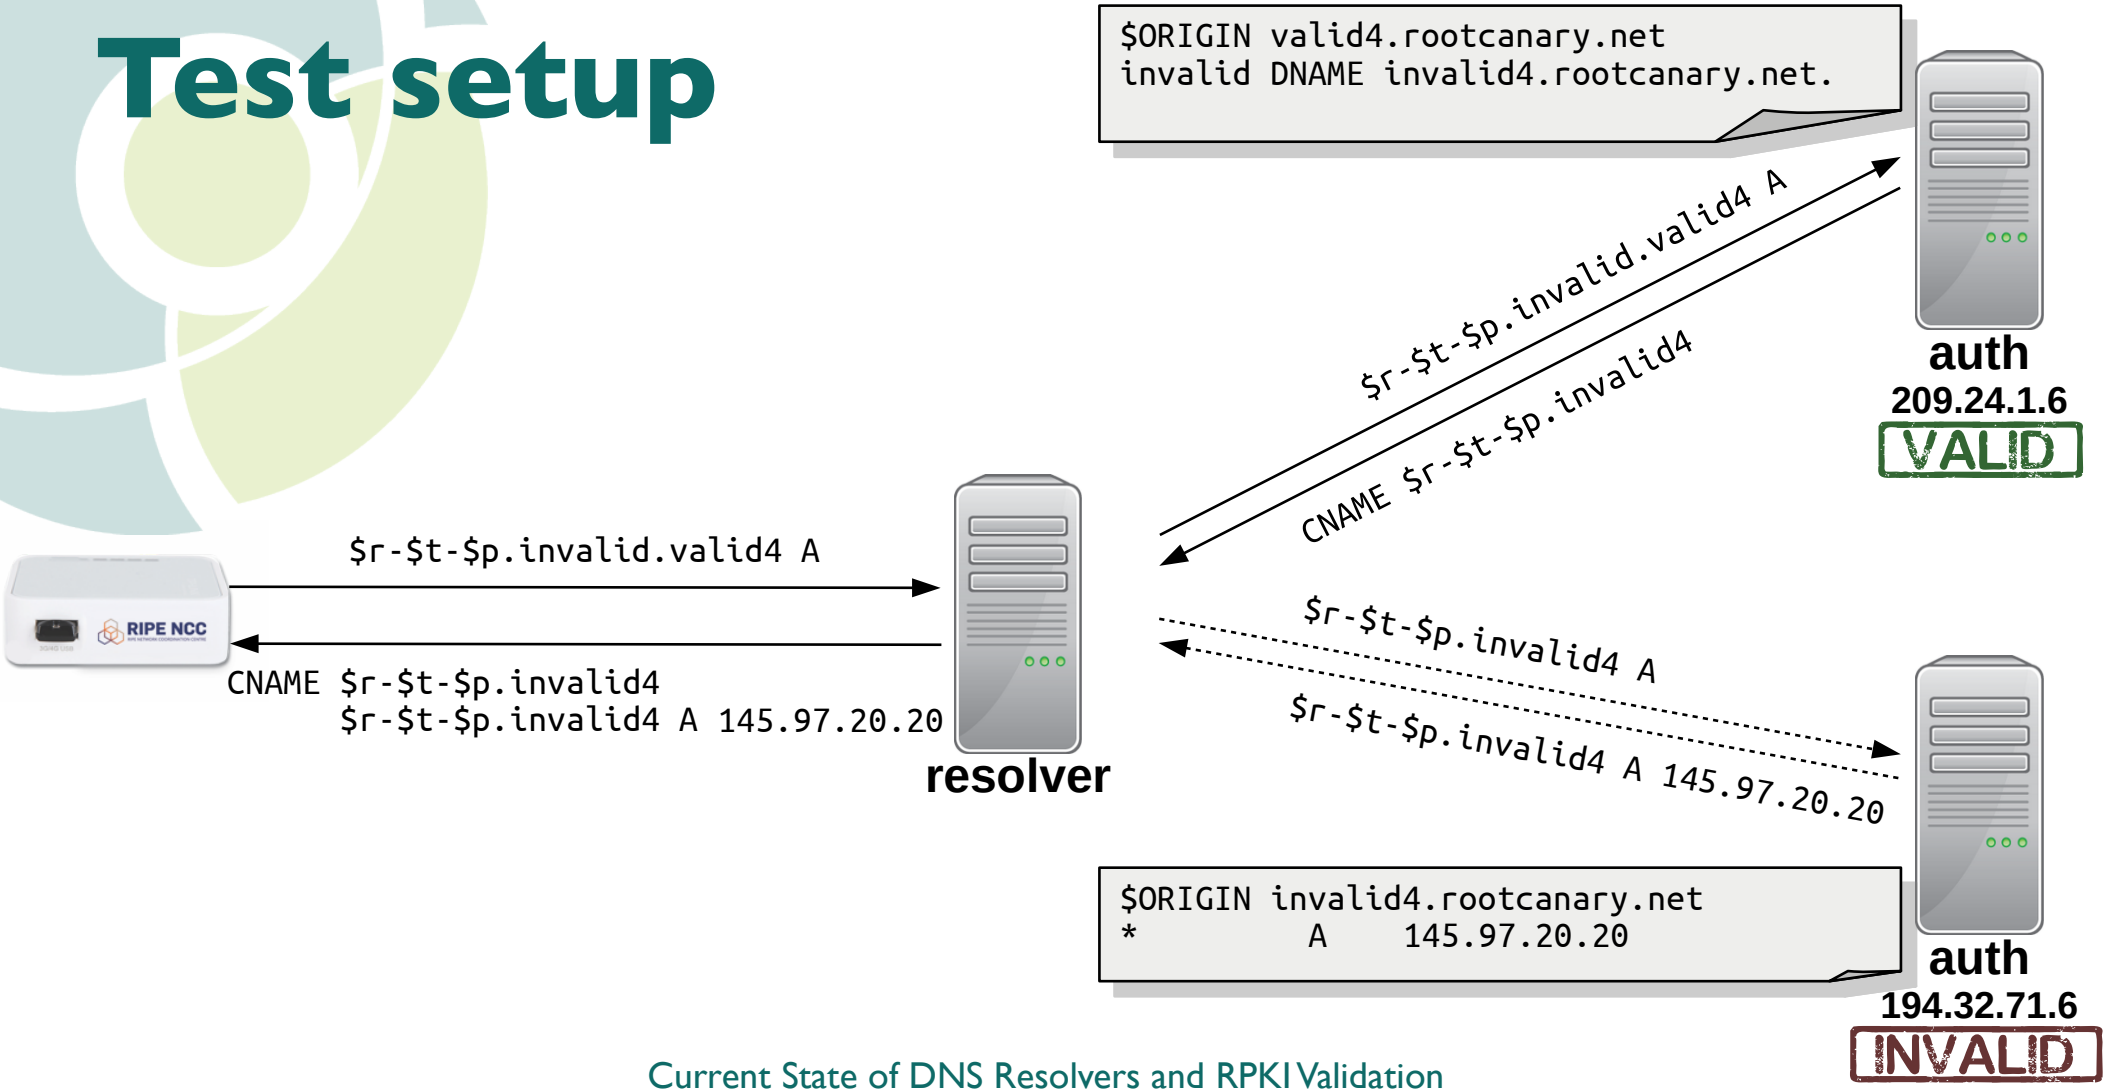
$ORIGIN valid4.rootcanary.net
invalid DNAME invalid4.rootcanary.net.
auth
209.24.1.6
VALID
$r-$t-$p.invalid.valid4 A
CNAME $r-$t-$p.invalid4
$r-$t-$p.invalid4 A
CNAME $r-$t-$p.invalid4
$r-$t-$p.invalid4 A 145.97.20.20
RIPE NCC
resolver
$ORIGIN invalid4.rootcanary.net
* A 145.97.20.20
auth
194.32.71.6
INVALID

Test setup

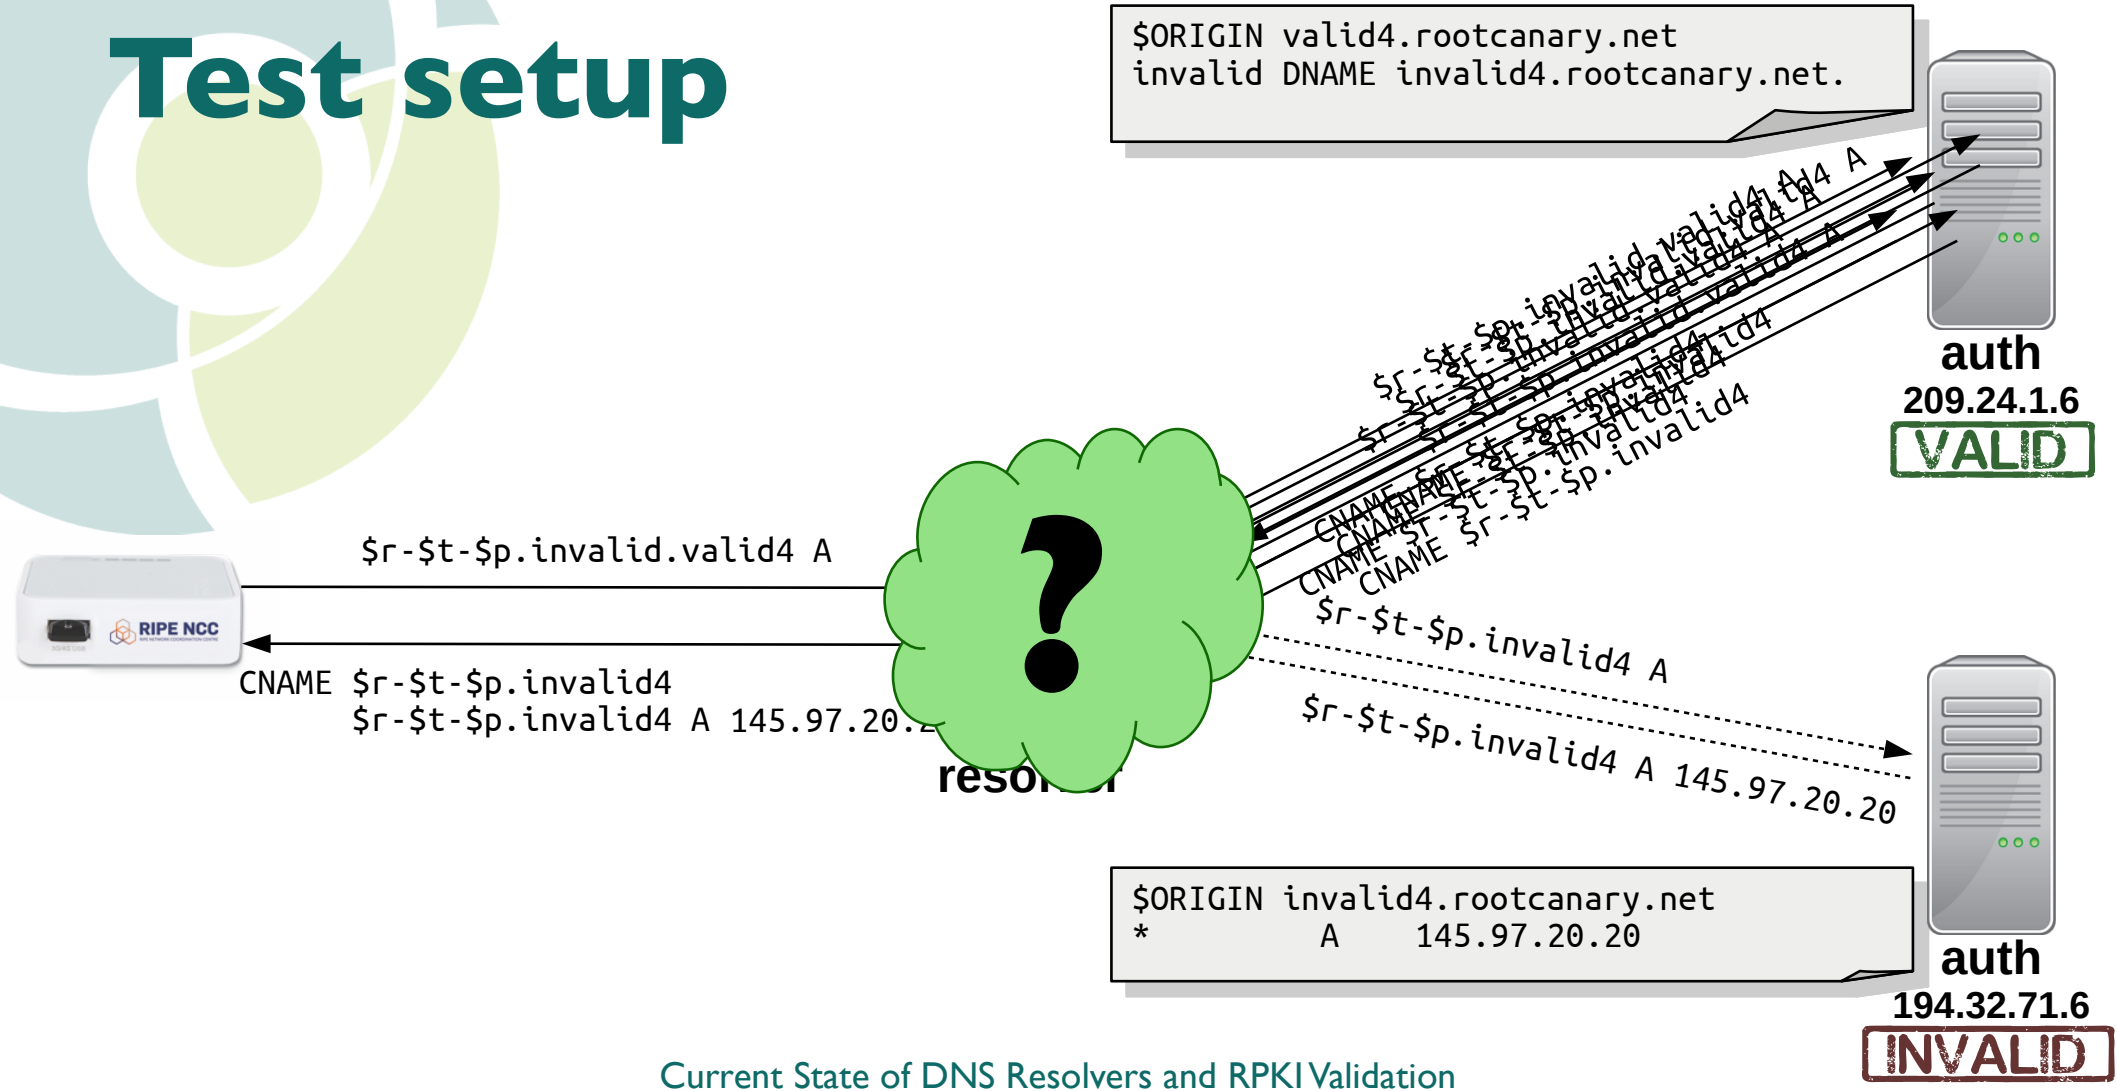
$ORIGIN valid4.rootcanary.net
invalid DNAME invalid4.rootcanary.net.
auth
209.24.1.6
VALID
$r-$t-$p.invalid.valid4 A
CNAME $r-$t-$p.invalid4
$r-$t-$p.invalid4 A 145.97.20.2
?
$r-$t-$p.invalid4 A
$r-$t-$p.invalid4 A 145.97.20.20
$ORIGIN invalid4.rootcanary.net
* A 145.97.20.20
auth
194.32.71.6
INVALID

Test setup

- Atlas measurement
  kindly provided by Emile Aben
- Beacon for the authoritatives
  kindly provided by Job Snijders

Results

Probe/resolver pair

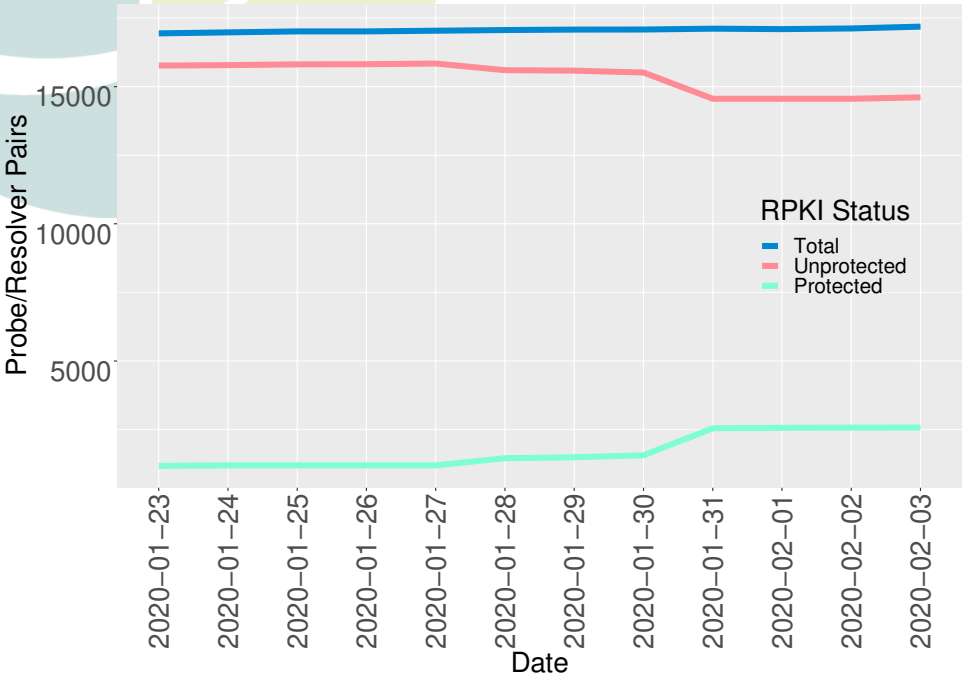

Probe time series

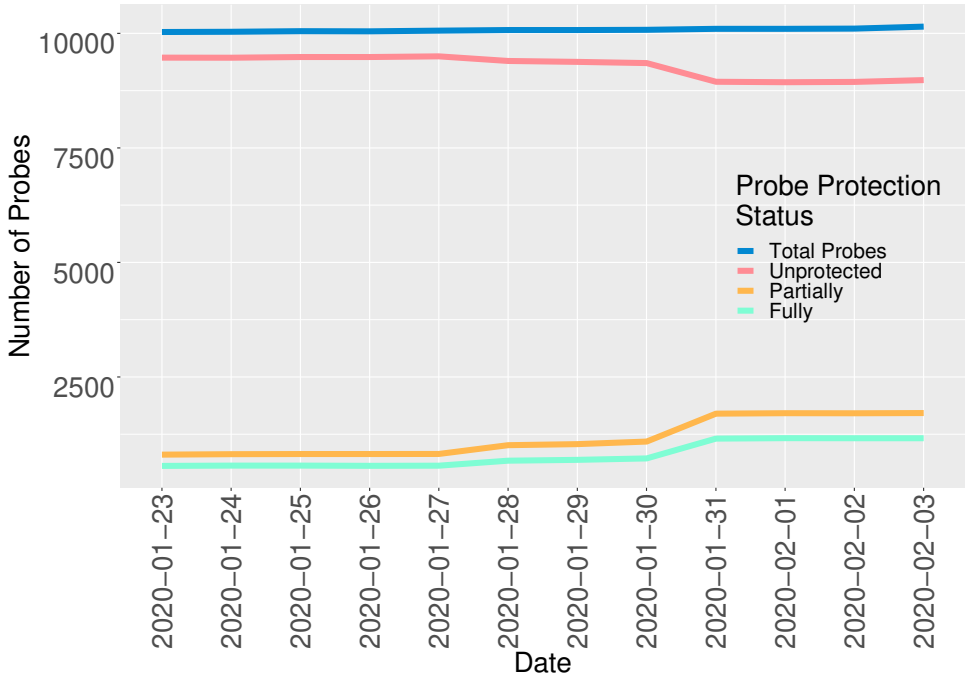

Results

Top ten most popular ASes

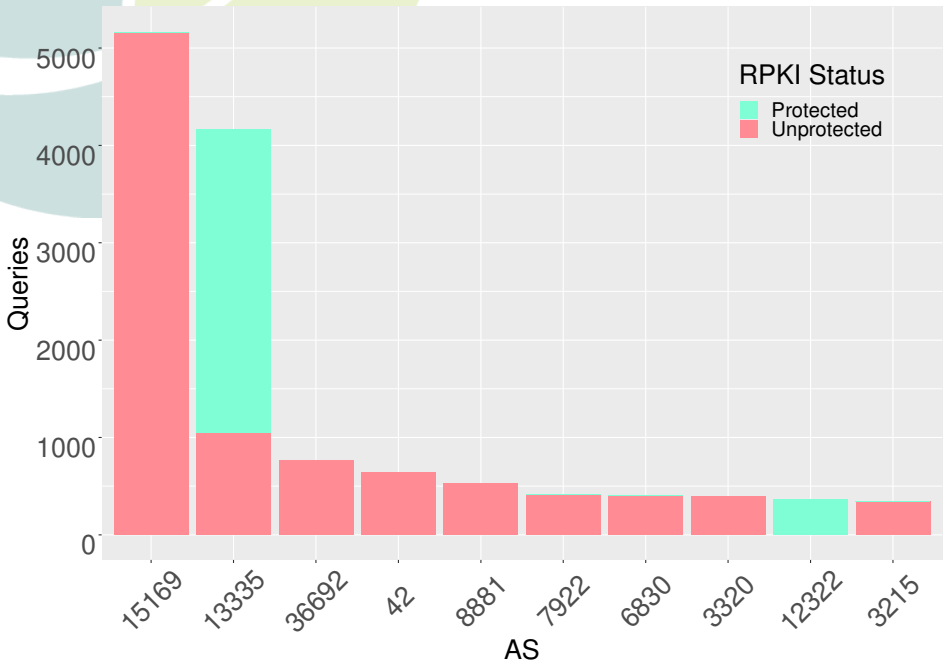

Top ten most protected ASes

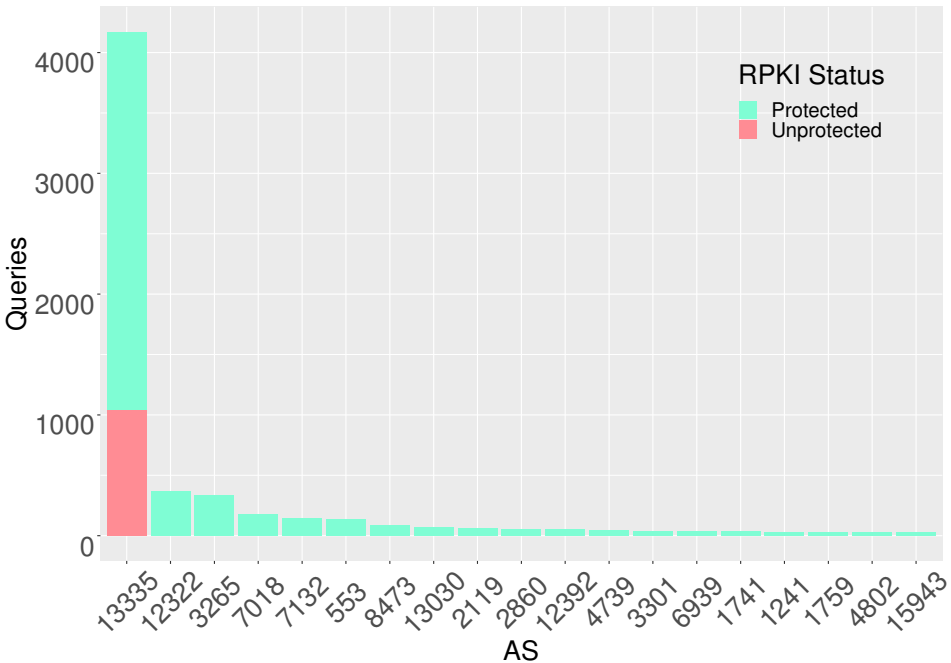

Results

Sub RQ: Does the length of the AS path matter?

Relationship RPKI protection and AS path length

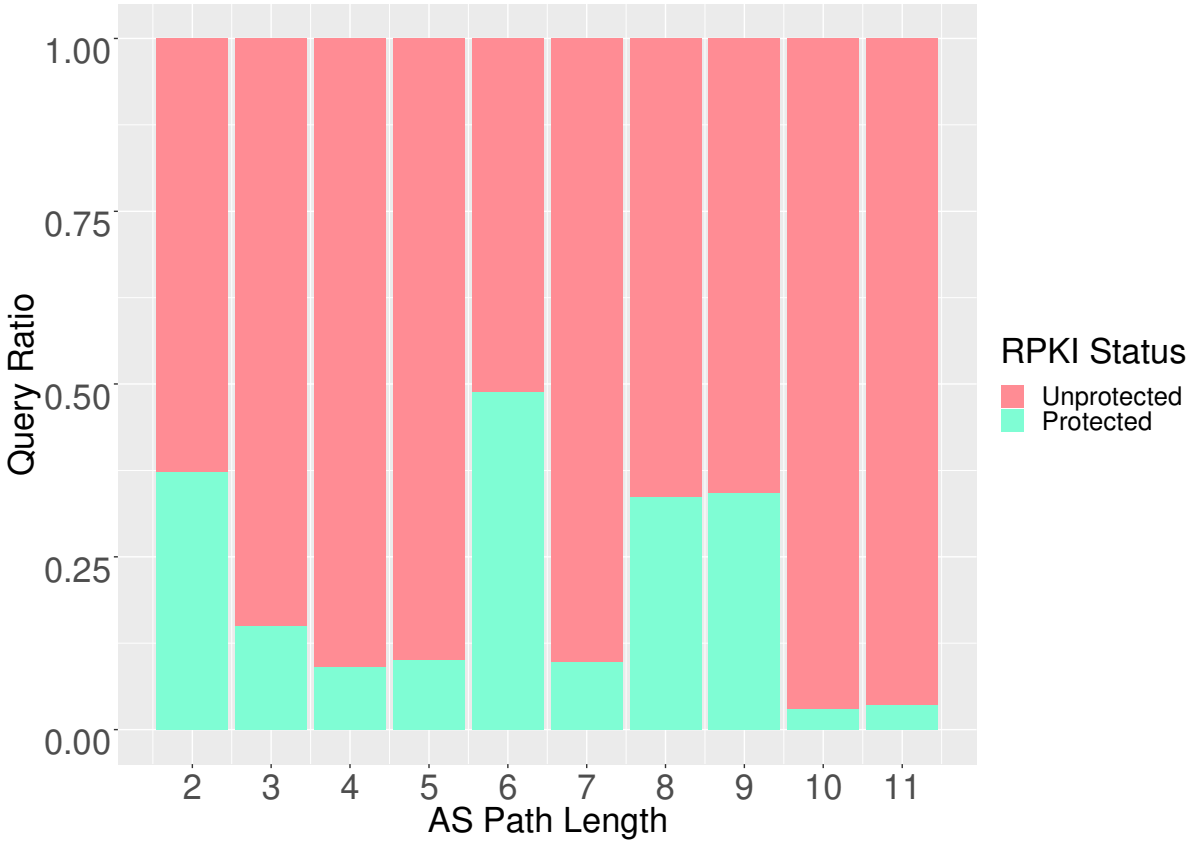

Results

Sub RQ: How does anycast influence protection?

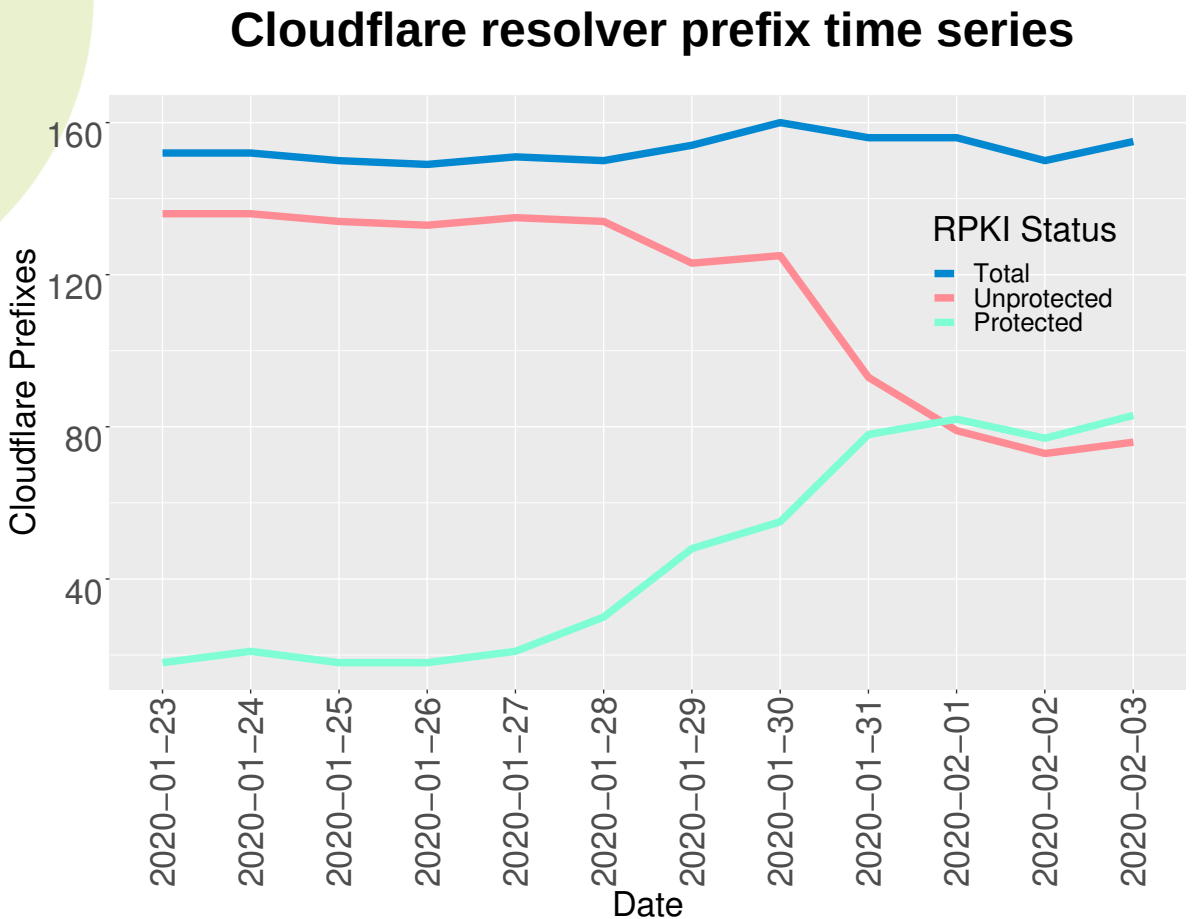

Current situation / IPv6

DNSThought

DNSThought

https://dnsthought.nlnetlabs.nl/does_rov4/#top_auth_asns

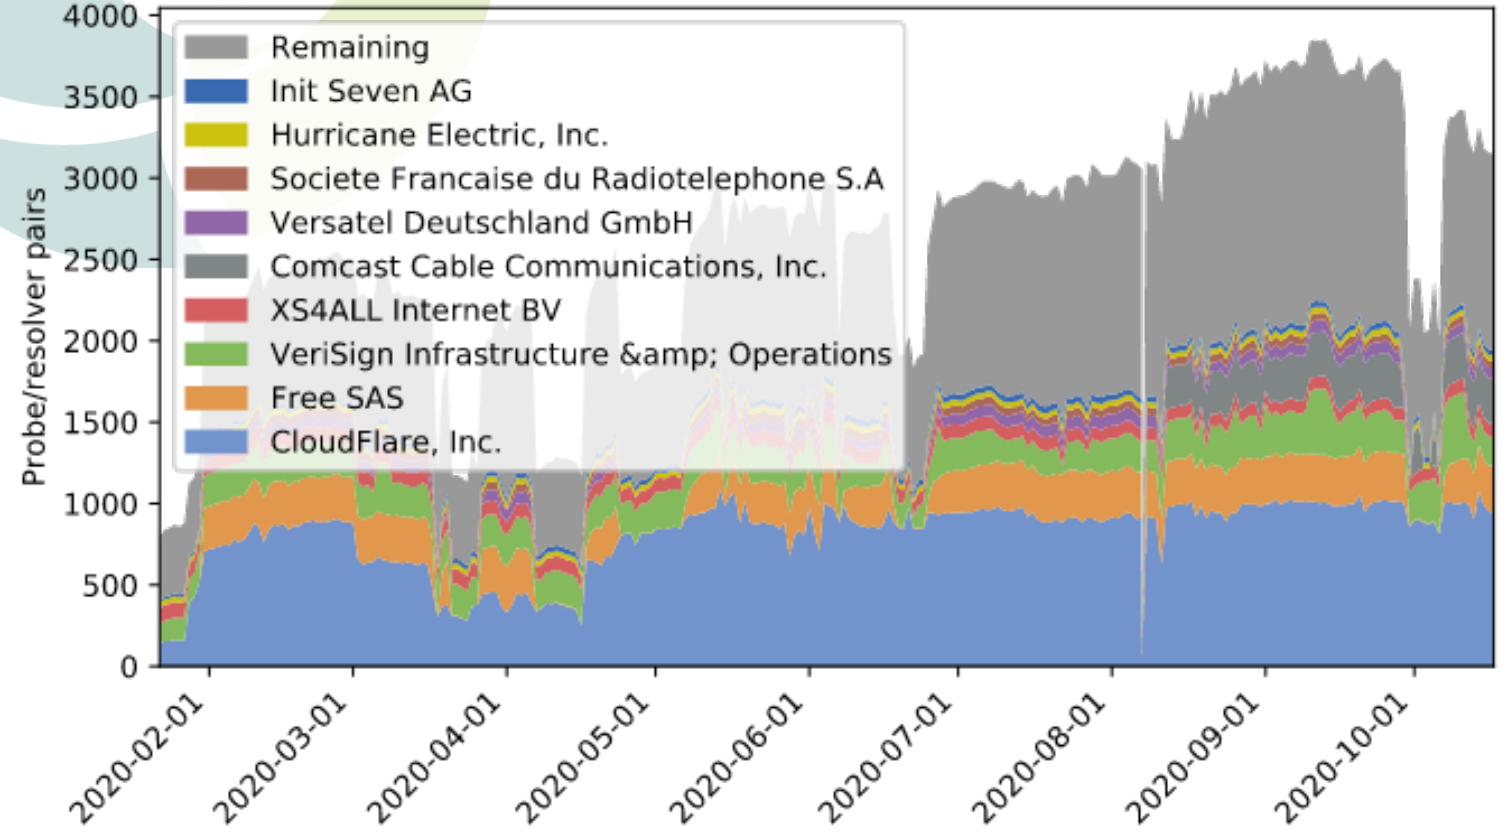

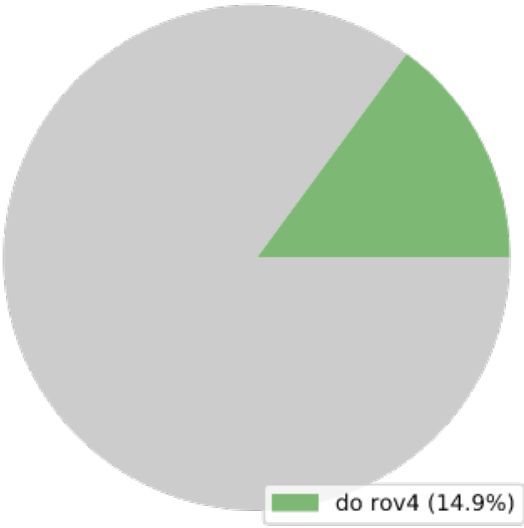

Test setup

```
$ORIGIN rootcanary.net
$TTL 60
@    SOA  ns1.surfnet.nl. (
          dns-beheer.surfnet.nl.
          2020080503 ; serial
          10800      ; refresh
          3600       ; retry
          604800     ; expire
          86400      ; minimum
          )
     NS   ns1.surfnet.nl.
     NS   ns2.surfnet.nl.
     NS   ns3.surfnet.nl.
     NS   ns1.zurich.surf.net.

$TTL 25200

valid6     NS   valid6
valid6     AAAA 2001:728:1808:5::6

invalid6   NS   invalid6
invalid6   AAAA 2001:7fb:fd04::6
```

```
$ORIGIN valid6.rootcanary.net
$TTL 300
@          SOA  valid6.rootcanary.net. (
                sysadm.rootcanary.org.
                2020012100 10800 3600
                604800 300 )
           NS   @
           A    2001:728:1808:5::6
$TTL 1
invalid    DNAME invalid6.rootcanary.net.
```

prefix	2001:728:1808::/48
max len	64
ASN	15562


```
$ORIGIN invalid6.rootcanary.net
$TTL 300
@          SOA  invalid6.rootcanary.net. (
                sysadm.rootcanary.org.
                2020012100 10800 3600
                604800 300 )
           NS   @
           A    2001:7fb:fd04::6
*          A    2001:610:188:408::20
```

prefix	2001:7fb:fd04::/48
max len	48
ASN	196615

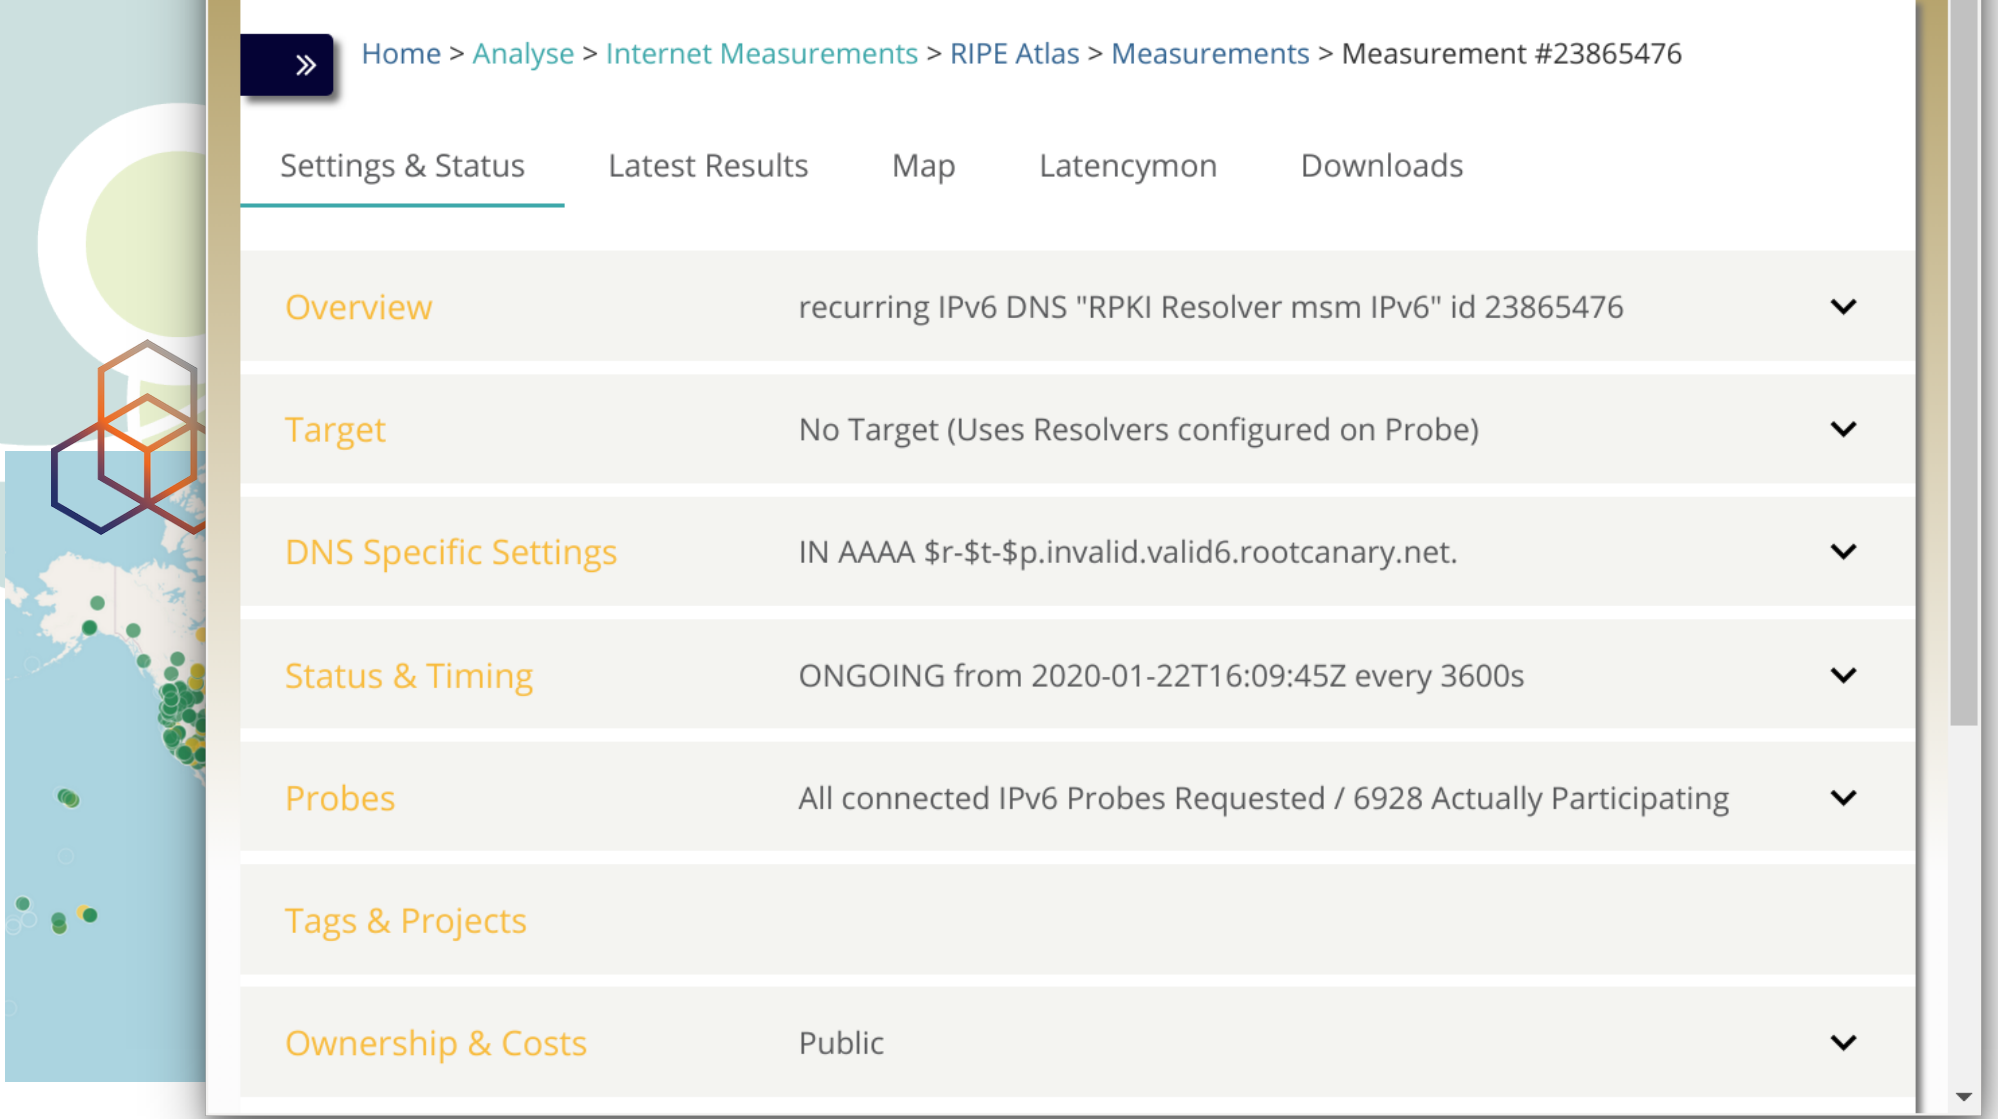
Home > Analyse > Internet Measurements > RIPE Atlas > Measurements > Measurement #23865476
Settings & Status
Latest Results
Map
Latencymon
Downloads
Overview
recurring IPv6 DNS "RPKI Resolver msm IPv6" id 23865476
Target
No Target (Uses Resolvers configured on Probe)
DNS Specific Settings
IN AAAA $r-$t-$p.invalid.valid6.rootcanary.net.
Status & Timing
ONGOING from 2020-01-22T16:09:45Z every 3600s
Probes
All connected IPv6 Probes Requested / 6928 Actually Participating
Tags & Projects
Ownership & Costs
Public

DNSThought

https://dnsthought.nlnetlabs.nl/does_rov6/#top_auth_asns

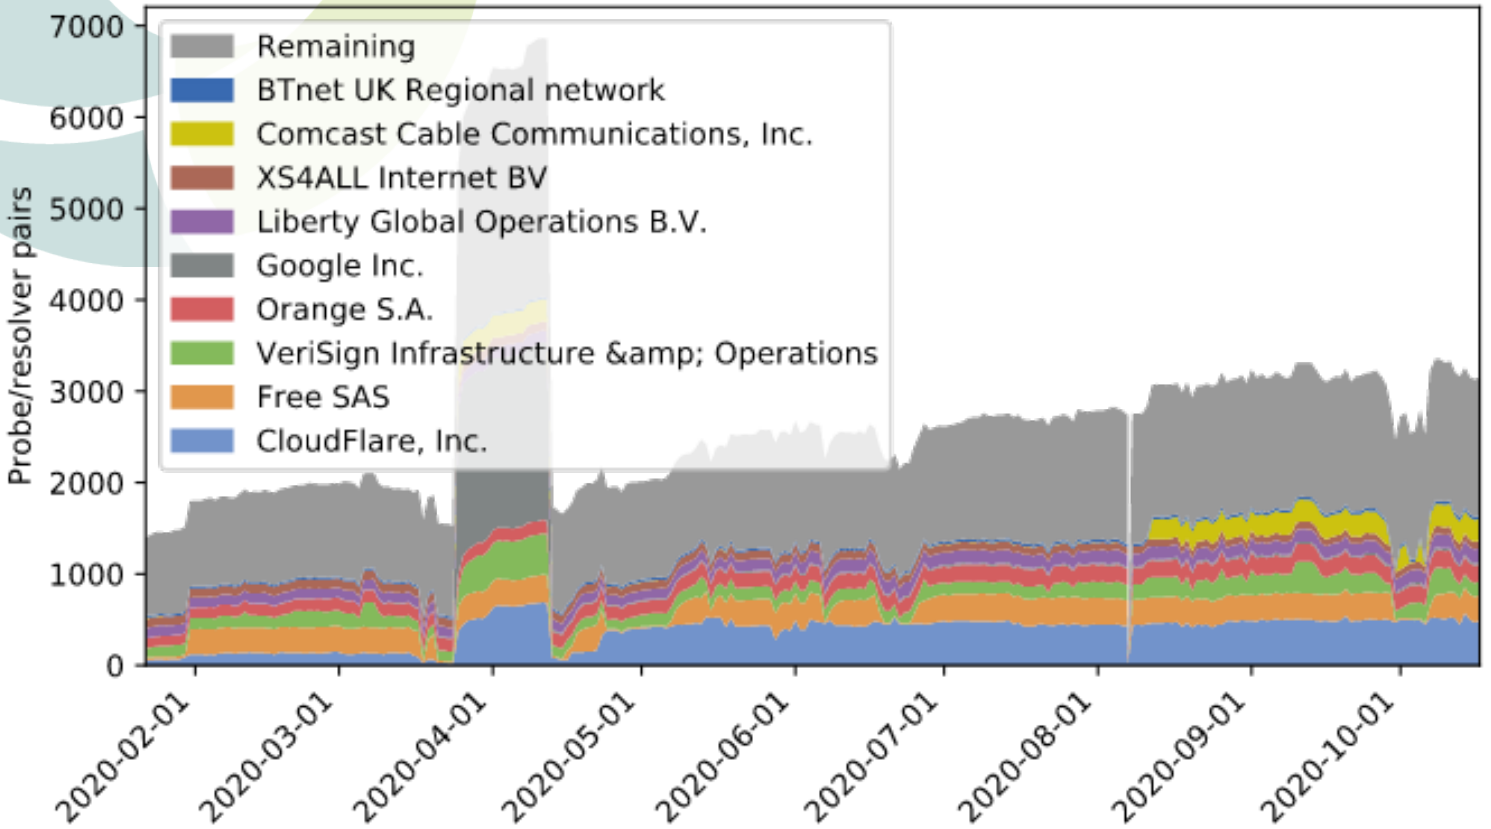

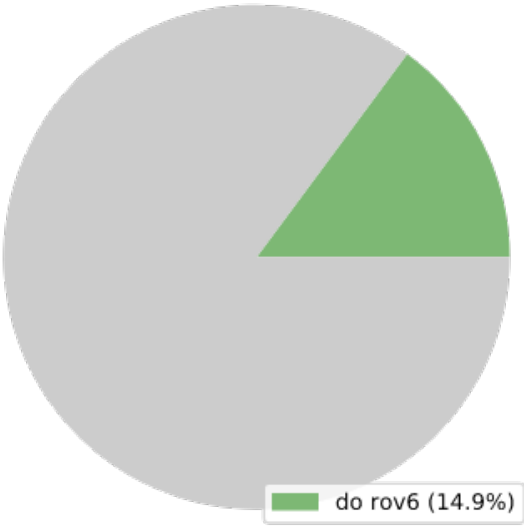

DNSThought

AS12322 Free SAS

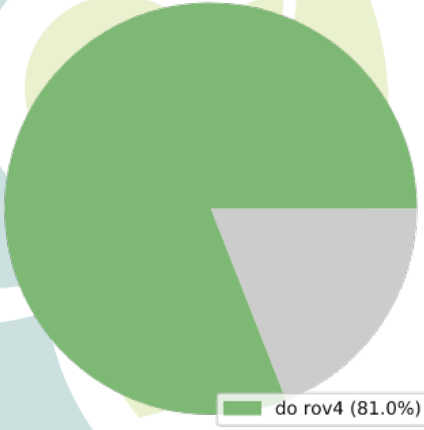

IPv4

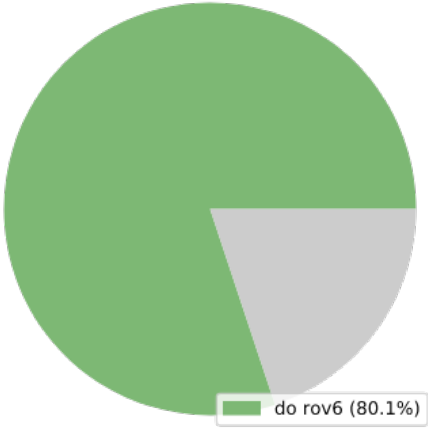

IPv6

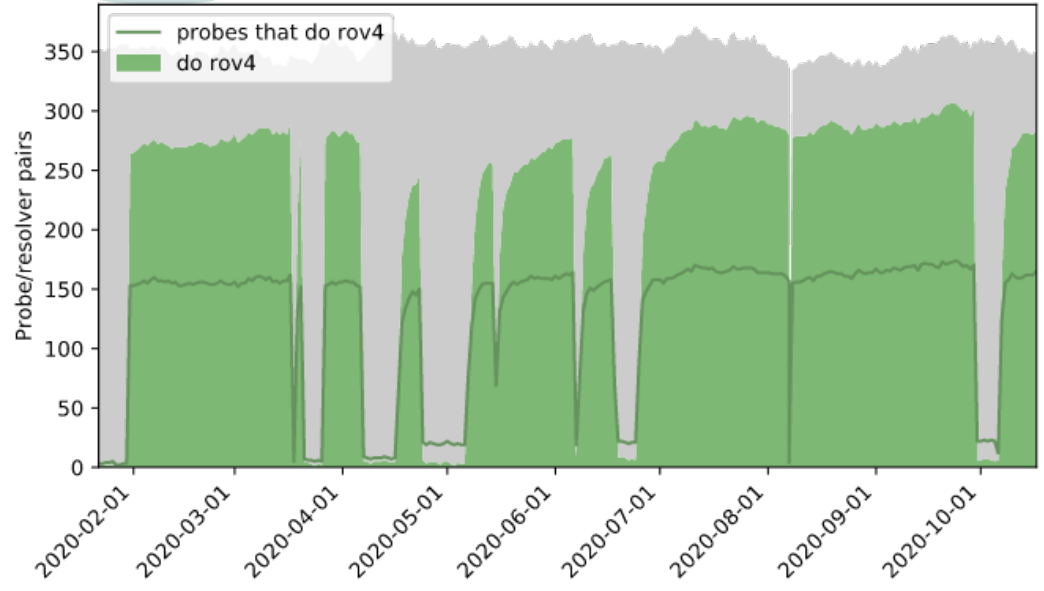

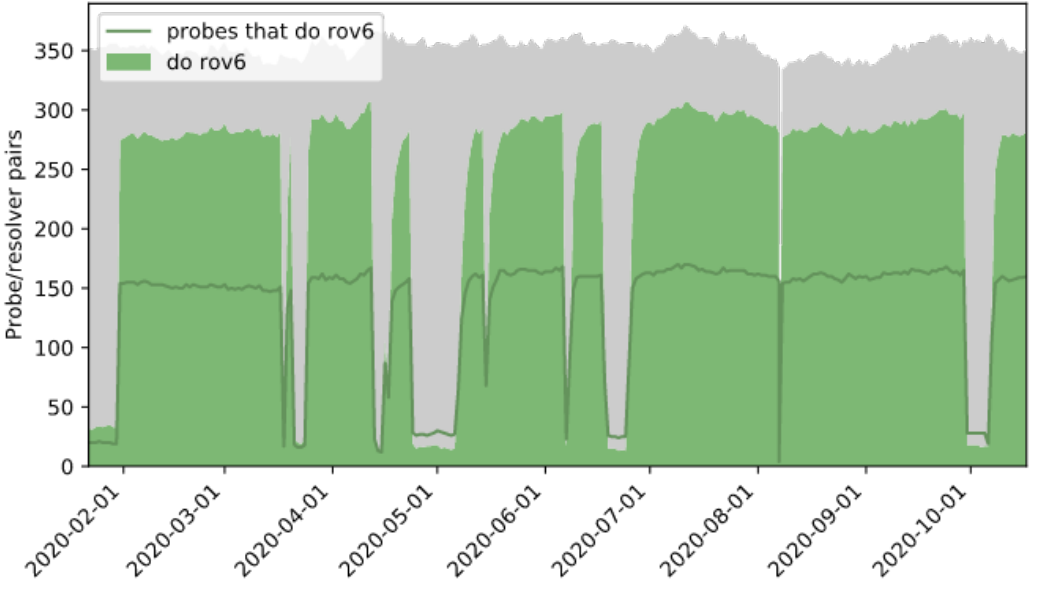

DNSThought

AS13335 Cloudflare

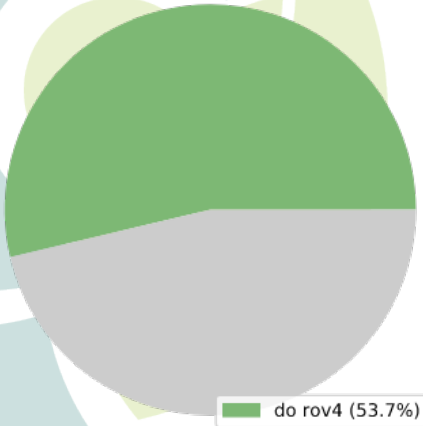

IPv4

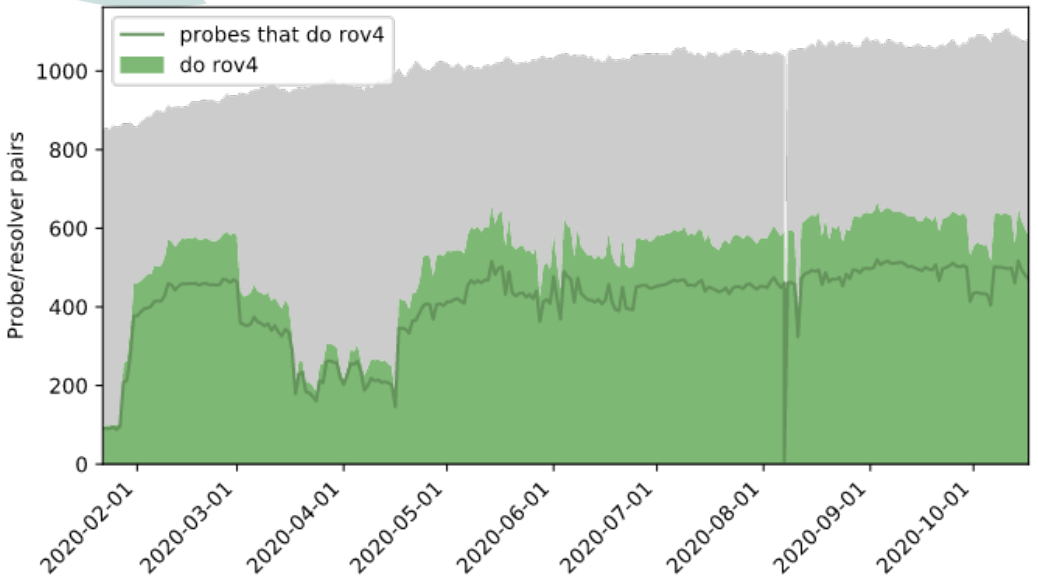

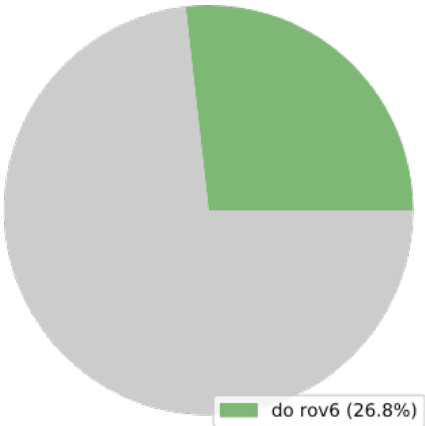

IPv6

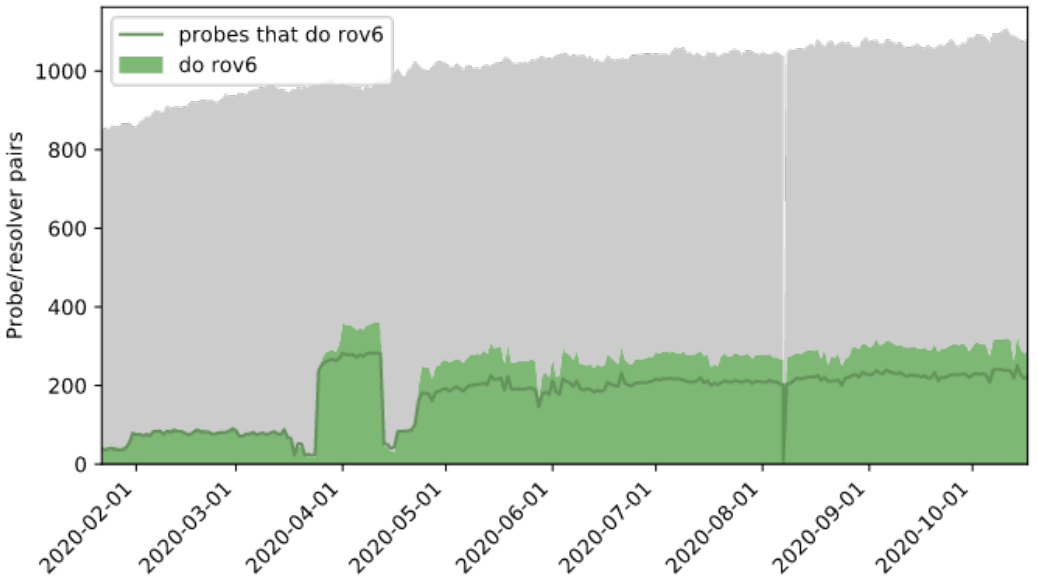

DNSThought

AS13335 Cloudflare

IPv6

Re: rov for cloudflare quad-1 resolver - Postvak IN - willem@nlnetlabs.nl - Mozill

Postvak IN - willem@nlnetlabs.

Re: rov for cloudflare quad-

Doorsturen | Archiveren | Ongewenst | Verwijderen | Meer

Van Louis Poinsignon <louis@cloudflare.com>

Onderwerp **Re: rov for cloudflare quad-1 resolver** 04-02-2020 16:39

Aan Martin J. Levy <martin@cloudflare.com>

Cc Emile Aben <emile.aben@ripe.net>, mij <willem@nlnetlabs

Once this is done, we should have all our routers dropping invalids.

I'm assuming you're running DNS tests through Atlas?
Could you run a TXT CH bind.hostname, it should return the PoP it's hitting.
This could help me trace any issues.

You have to know that we keep a default route.
This should avoid a hijack sent on a peering session for instance but a transit may still have the route in it's table.

Best,
Louis

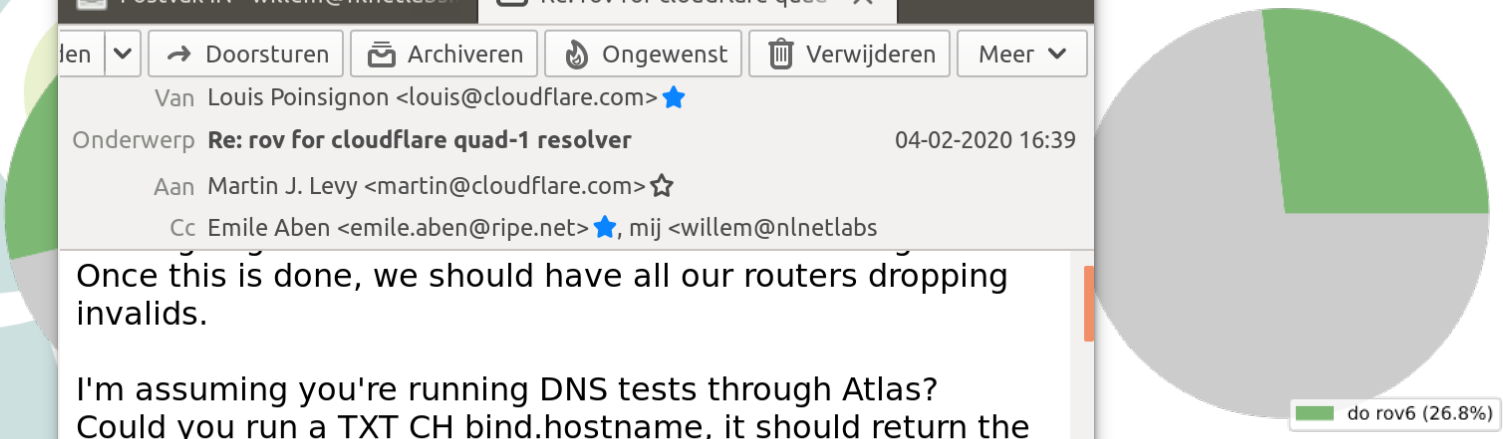

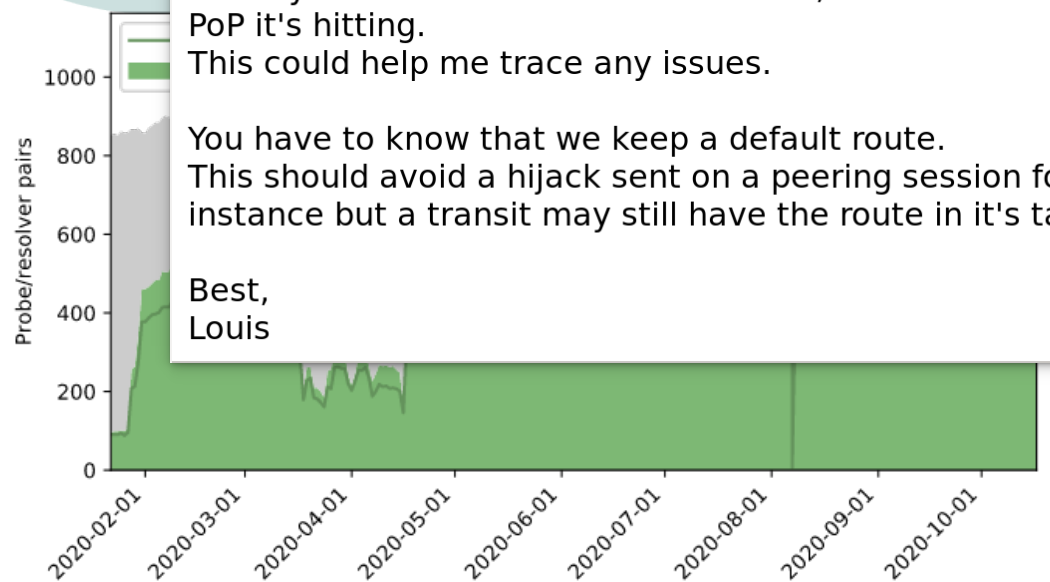

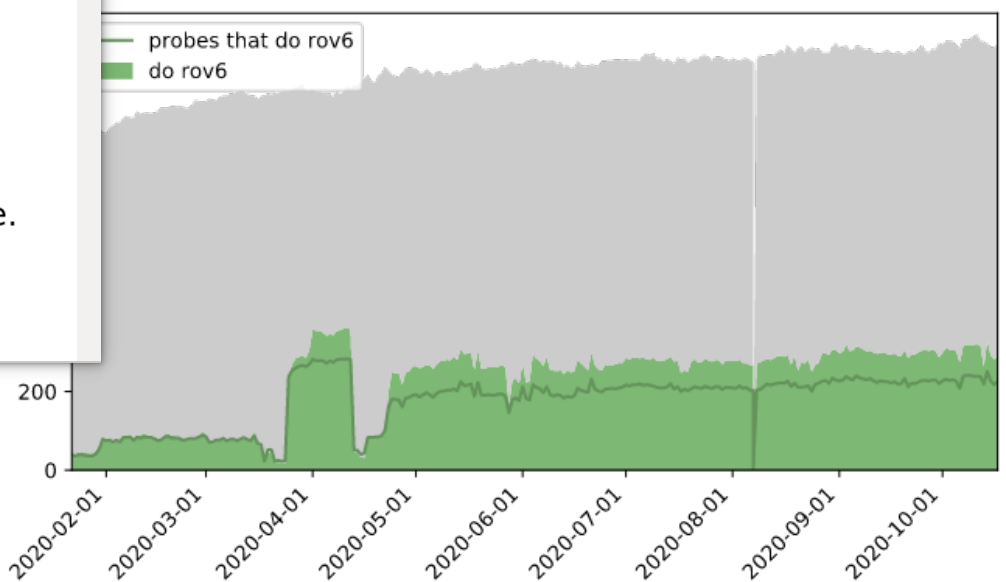

DNSThought

AS13335 Cloudflare

IPv6

Re: rov for cloudflare quad-1 resolver - Postvak IN - willem@nlnetlabs.nl - Mozill

Postvak IN - willem@nlnetlabs.

Re: rov for cloudflare quad-

Doorsturen | Archiveren | Ongewenst | Verwijderen | Meer

Van Louis Poinsignon <louis@cloudflare.com>

Onderwerp

Aan

Once t
invalids

I'm ass
Could y
PoP it's
This co

You hav
This sh
instanc

Best,
Louis

Re: rov for cloudflare quad-1 resolver - Postvak IN - willem@nlnetlabs.nl - Mozilla Thunde

Postvak IN - willem@nlnetlabs.

Re: rov for cloudflare quad-

Beantwoorden | Doorsturen | Archiveren | Ongewenst | Verwijderen | Meer

Van Martin J. Levy <martin@cloudflare.com>

Onderwerp **Re: rov for cloudflare quad-1 resolver** 04-02-2020 16:44

Aan Louis Poinsignon <louis@cloudflare.com>

Cc Emile Aben <emile.aben@ripe.net> mij <willem@nlnetlabs.nl>

Louis,

Let's make that: `TXT CH id.server` and match the newer RFCs. :)

Martin

> On Feb 4, 2020, at 7:40 AM, Louis Poinsignon
> <louis@cloudflare.com> wrote:

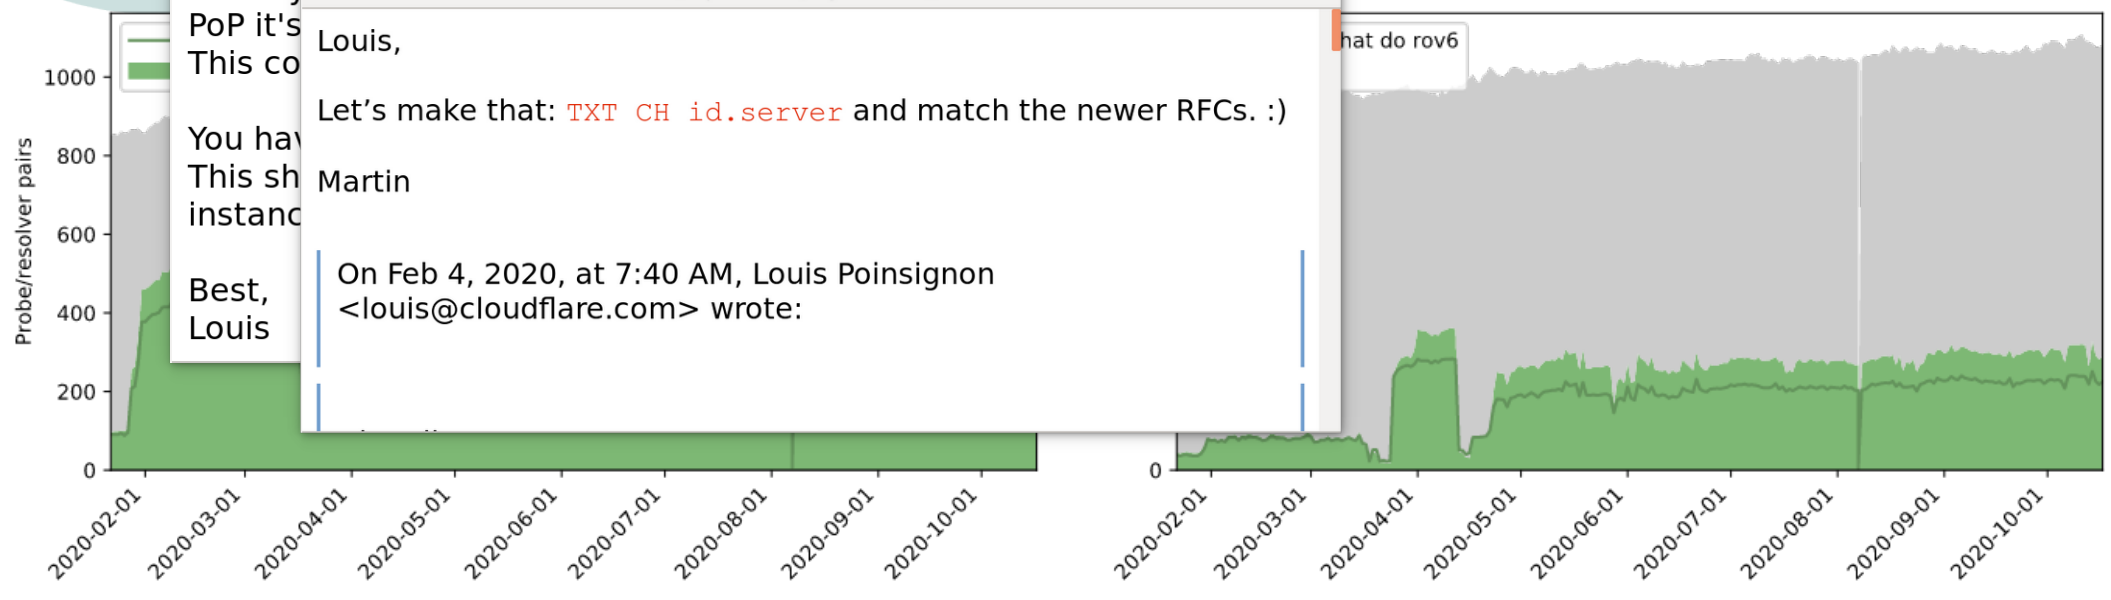

DNSThought

AS13335 Cloudflare

IPv6

Re: rov for cloudflare quad-1 resolver - Postvak IN - willem@nlnetlabs.nl - Mozill

Postvak IN - willem@nlnetlabs.

Re: rov for cloudflare quad-

Doorsturen | Archiveren | Ongewenst | Verwijderen | Meer

Van Louis Poinsignon <louis@cloudflare.com>

Onderwerp

Aan

Once t
invalids

I'm ass
Could y
PoP it's
This co

You hav
This sh
instanc

Best,
Louis

Re: rov for cloudflare quad-1 resolver - Postvak IN - willem@nlnetlabs.nl - Mozilla Thunde

Postvak IN - willem@nlnetlabs.

Re: rov for cloudflare quad-

Beantwoorden | Doorsturen | Archiveren | Ongewenst | Verwijderen | Meer

Van Martin J. Levy <martin@cloudflare.com>

Onderwerp **Re: rov for cloudflare quad-1 resolver**

Aan Louis Poinsignon <louis@cloudflare.com>

Cc Emile Aben <emile.aben@ripe.net> mij <willem

Louis,

Let's make that: `TXT CH id.server` and

Martin

> On Feb 4, 2020, at 7:40 AM, Louis Poinsi
> <louis@cloudflare.com> wrote:

willem@makaak: ~

```
willem@makaak:~$ dig @1.1.1.1 TXT CH id.server

; <<>> DiG 9.16.1-Ubuntu <<>> @1.1.1.1 TXT CH id.server
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 61529
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;id.server.                     CH      TXT

;; ANSWER SECTION:
id.server.              0       CH      TXT     "AMS"

;; Query time: 12 msec
;; SERVER: 1.1.1.1#53(1.1.1.1)
;; WHEN: zo okt 18 22:35:33 CEST 2020
;; MSG SIZE  rcvd: 43

willem@makaak:~$
```

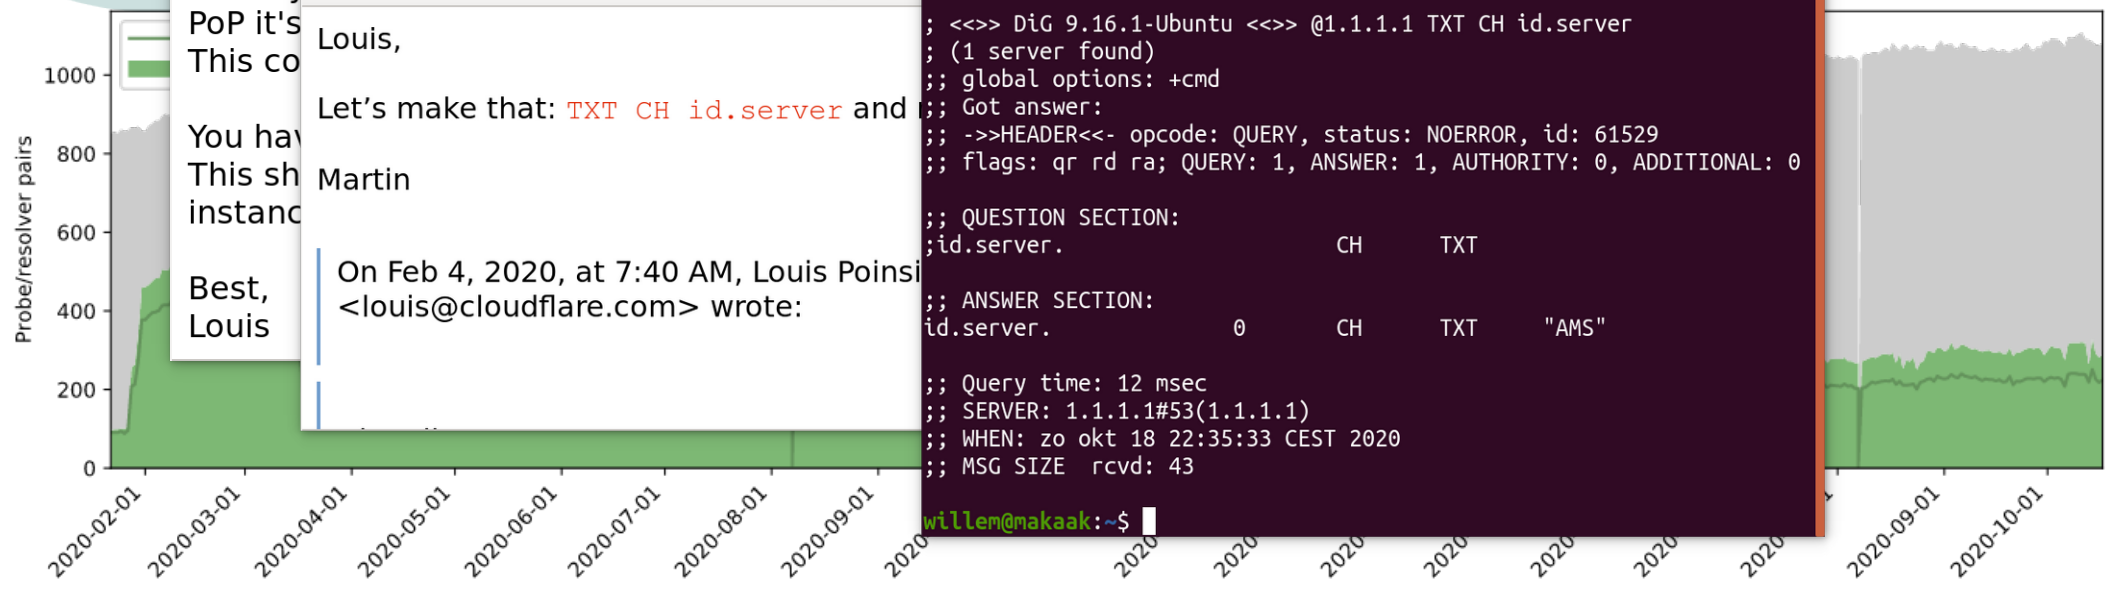

1.1.1.1 IPv4 RoV – PoPs

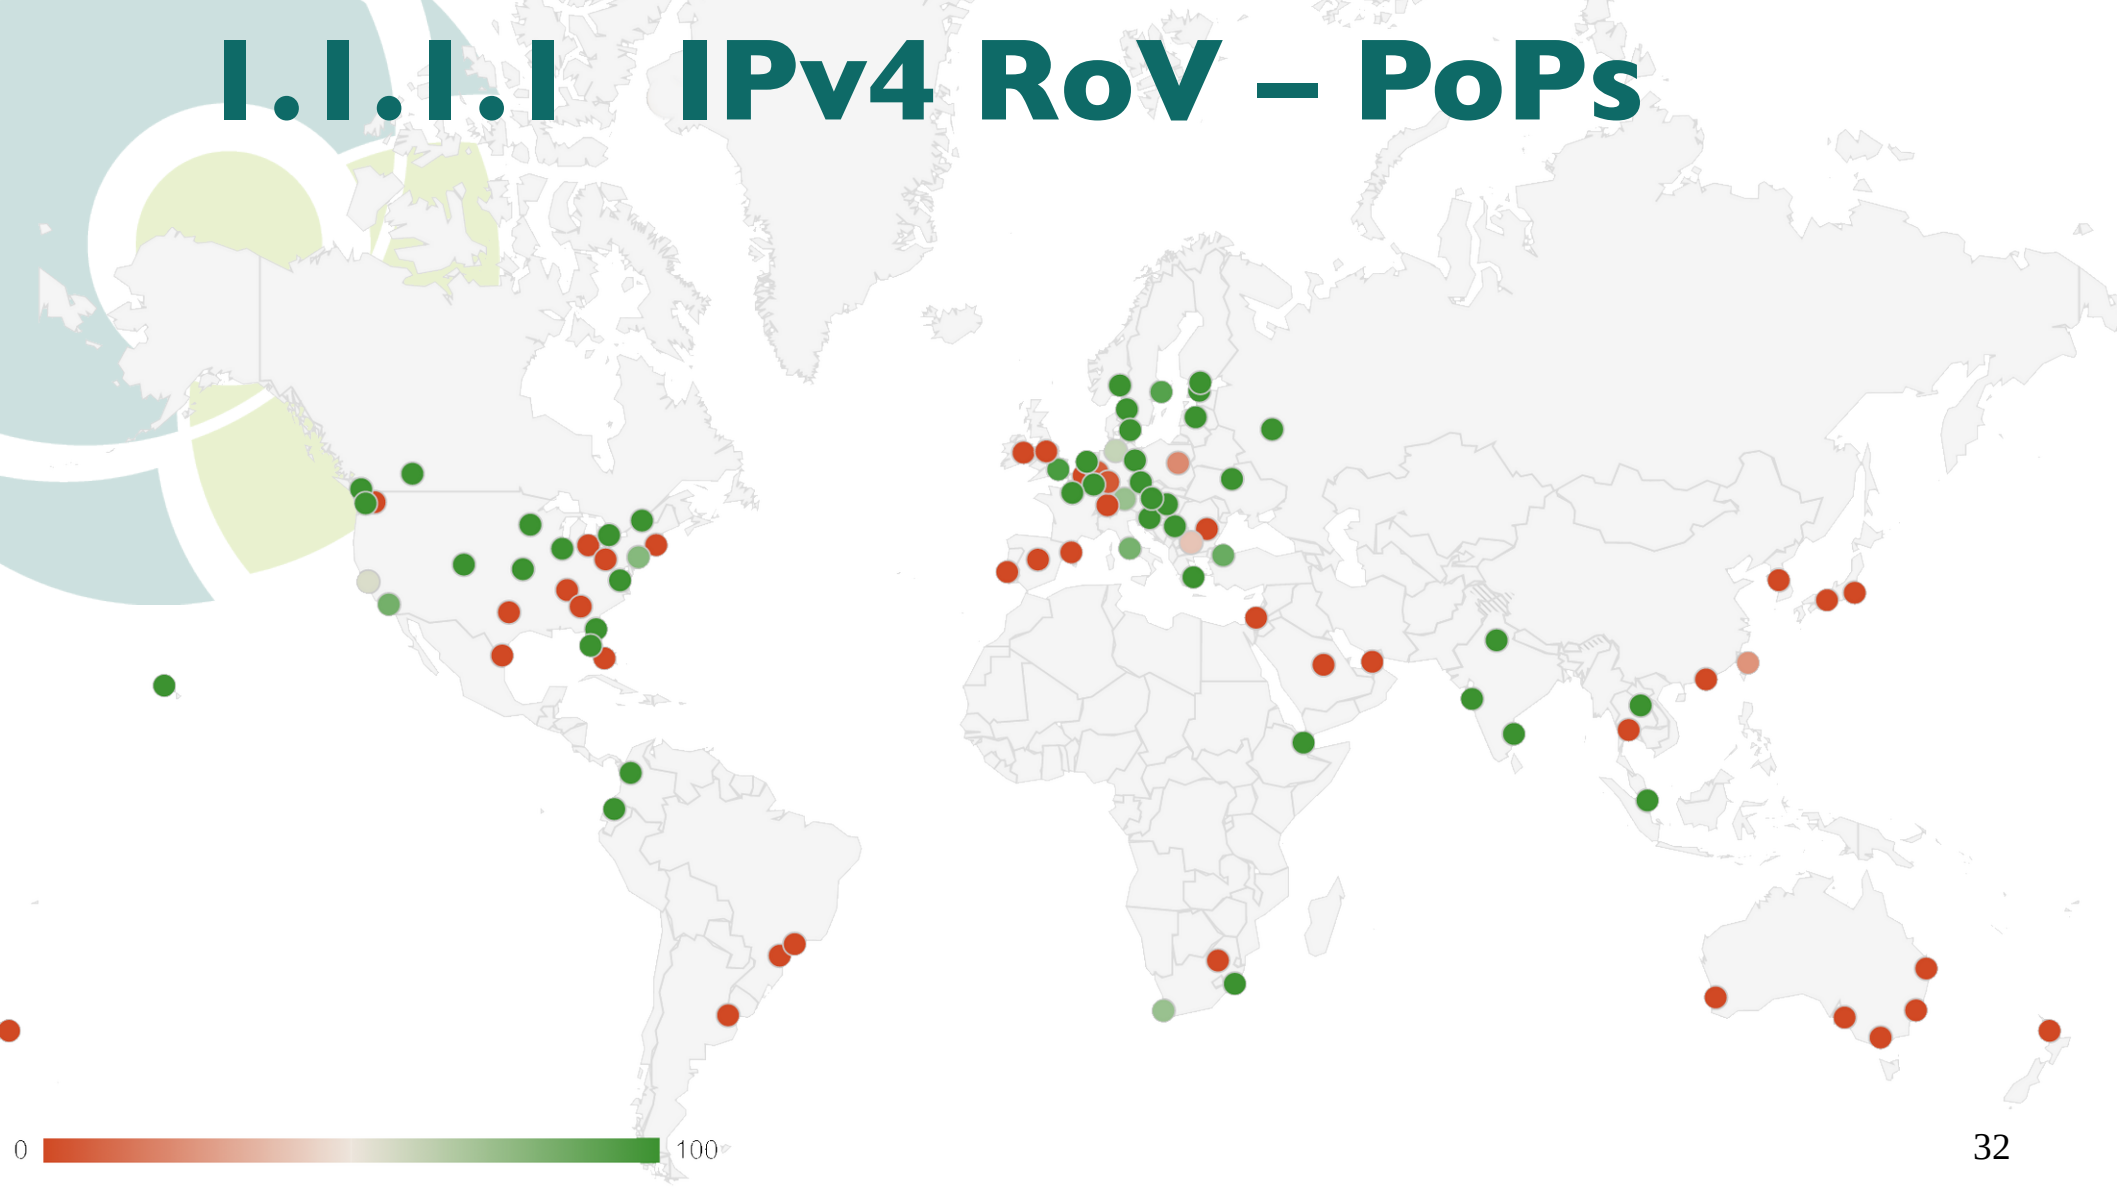

1.1.1.1 IPv4 RoV – Countries

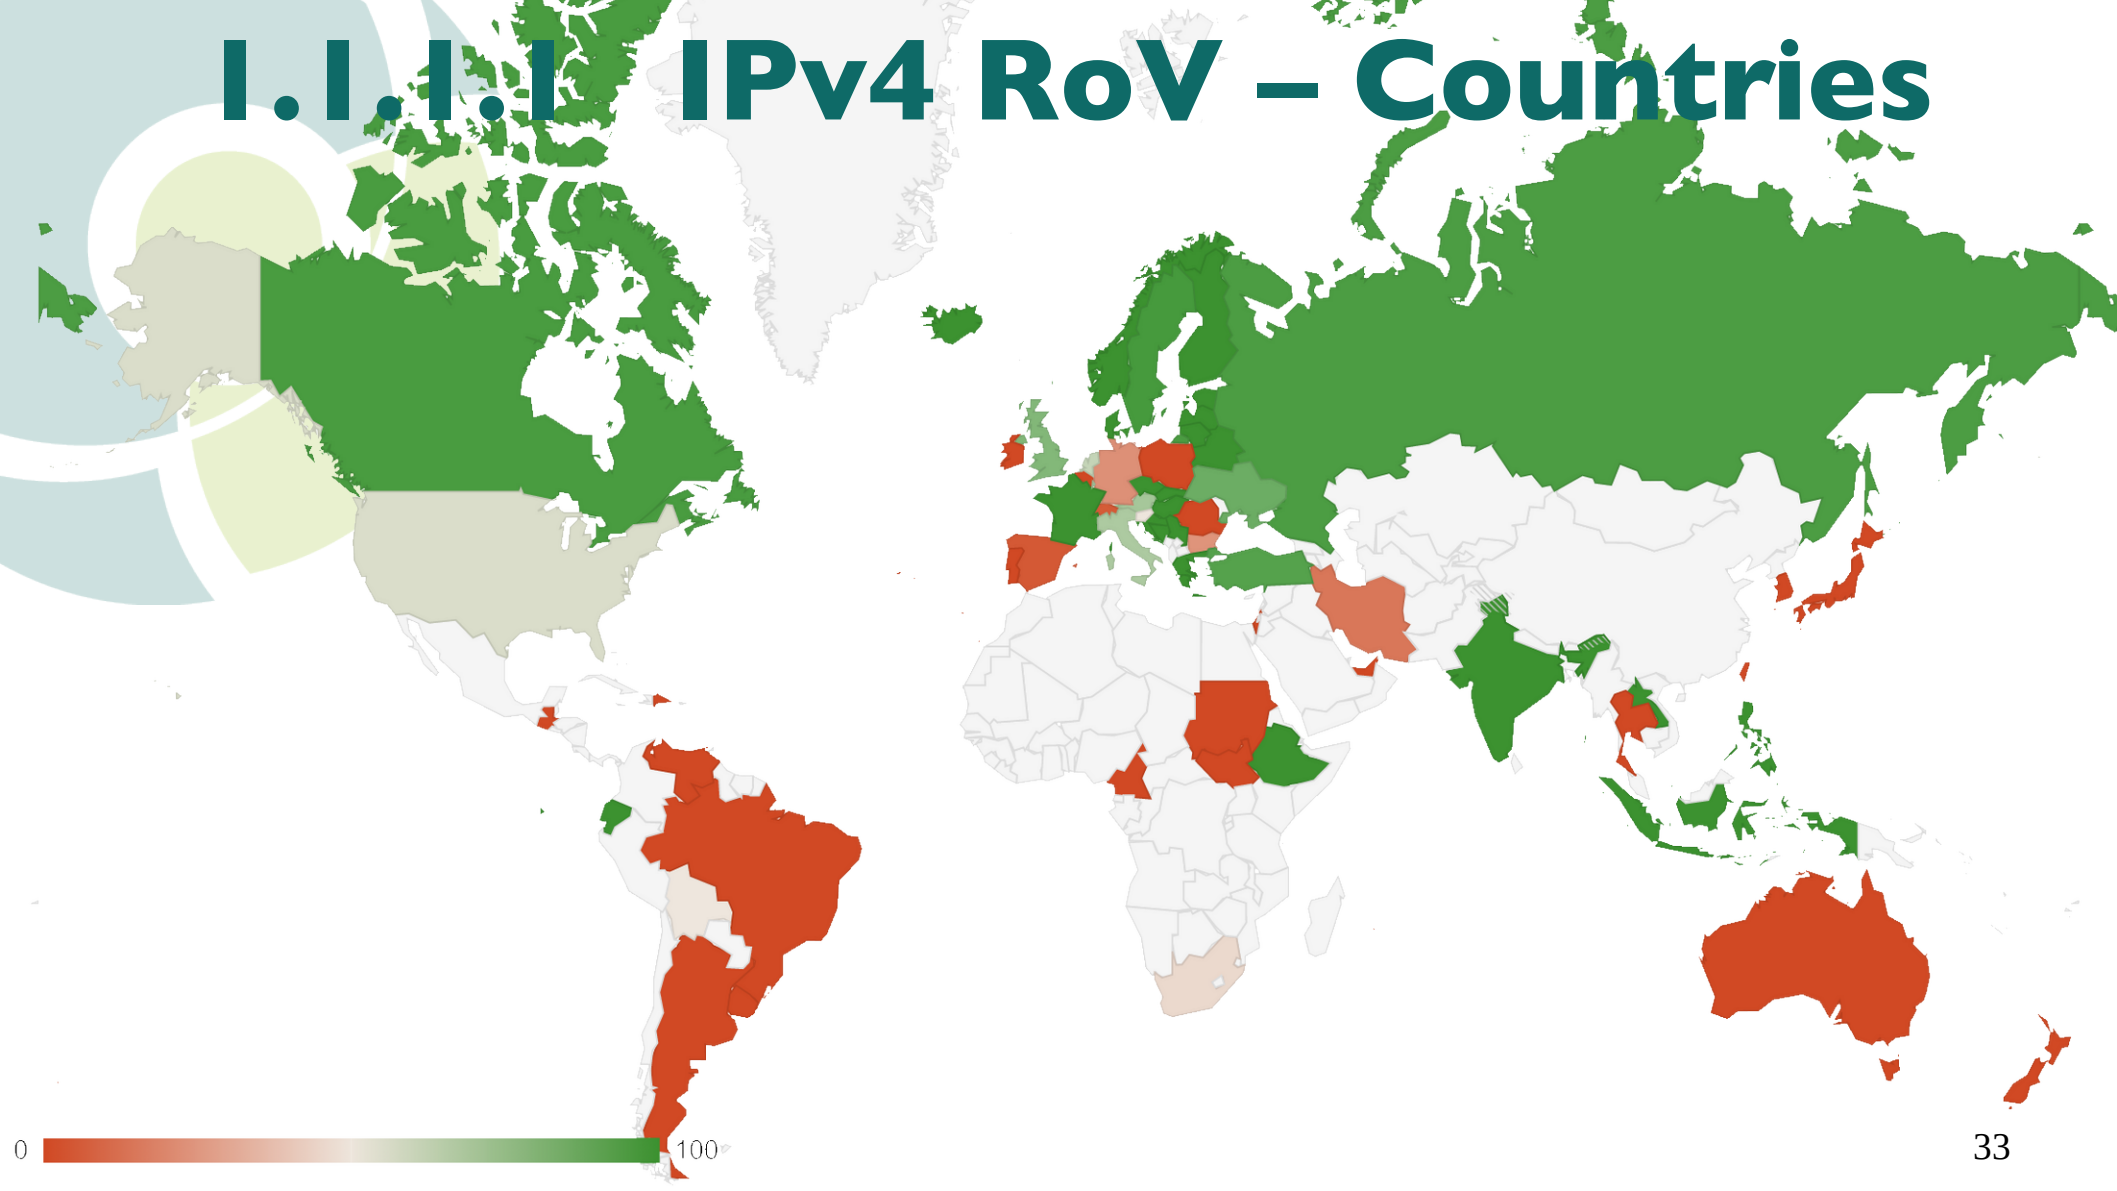

1.1.1.1 IPv6 RoV – PoPs

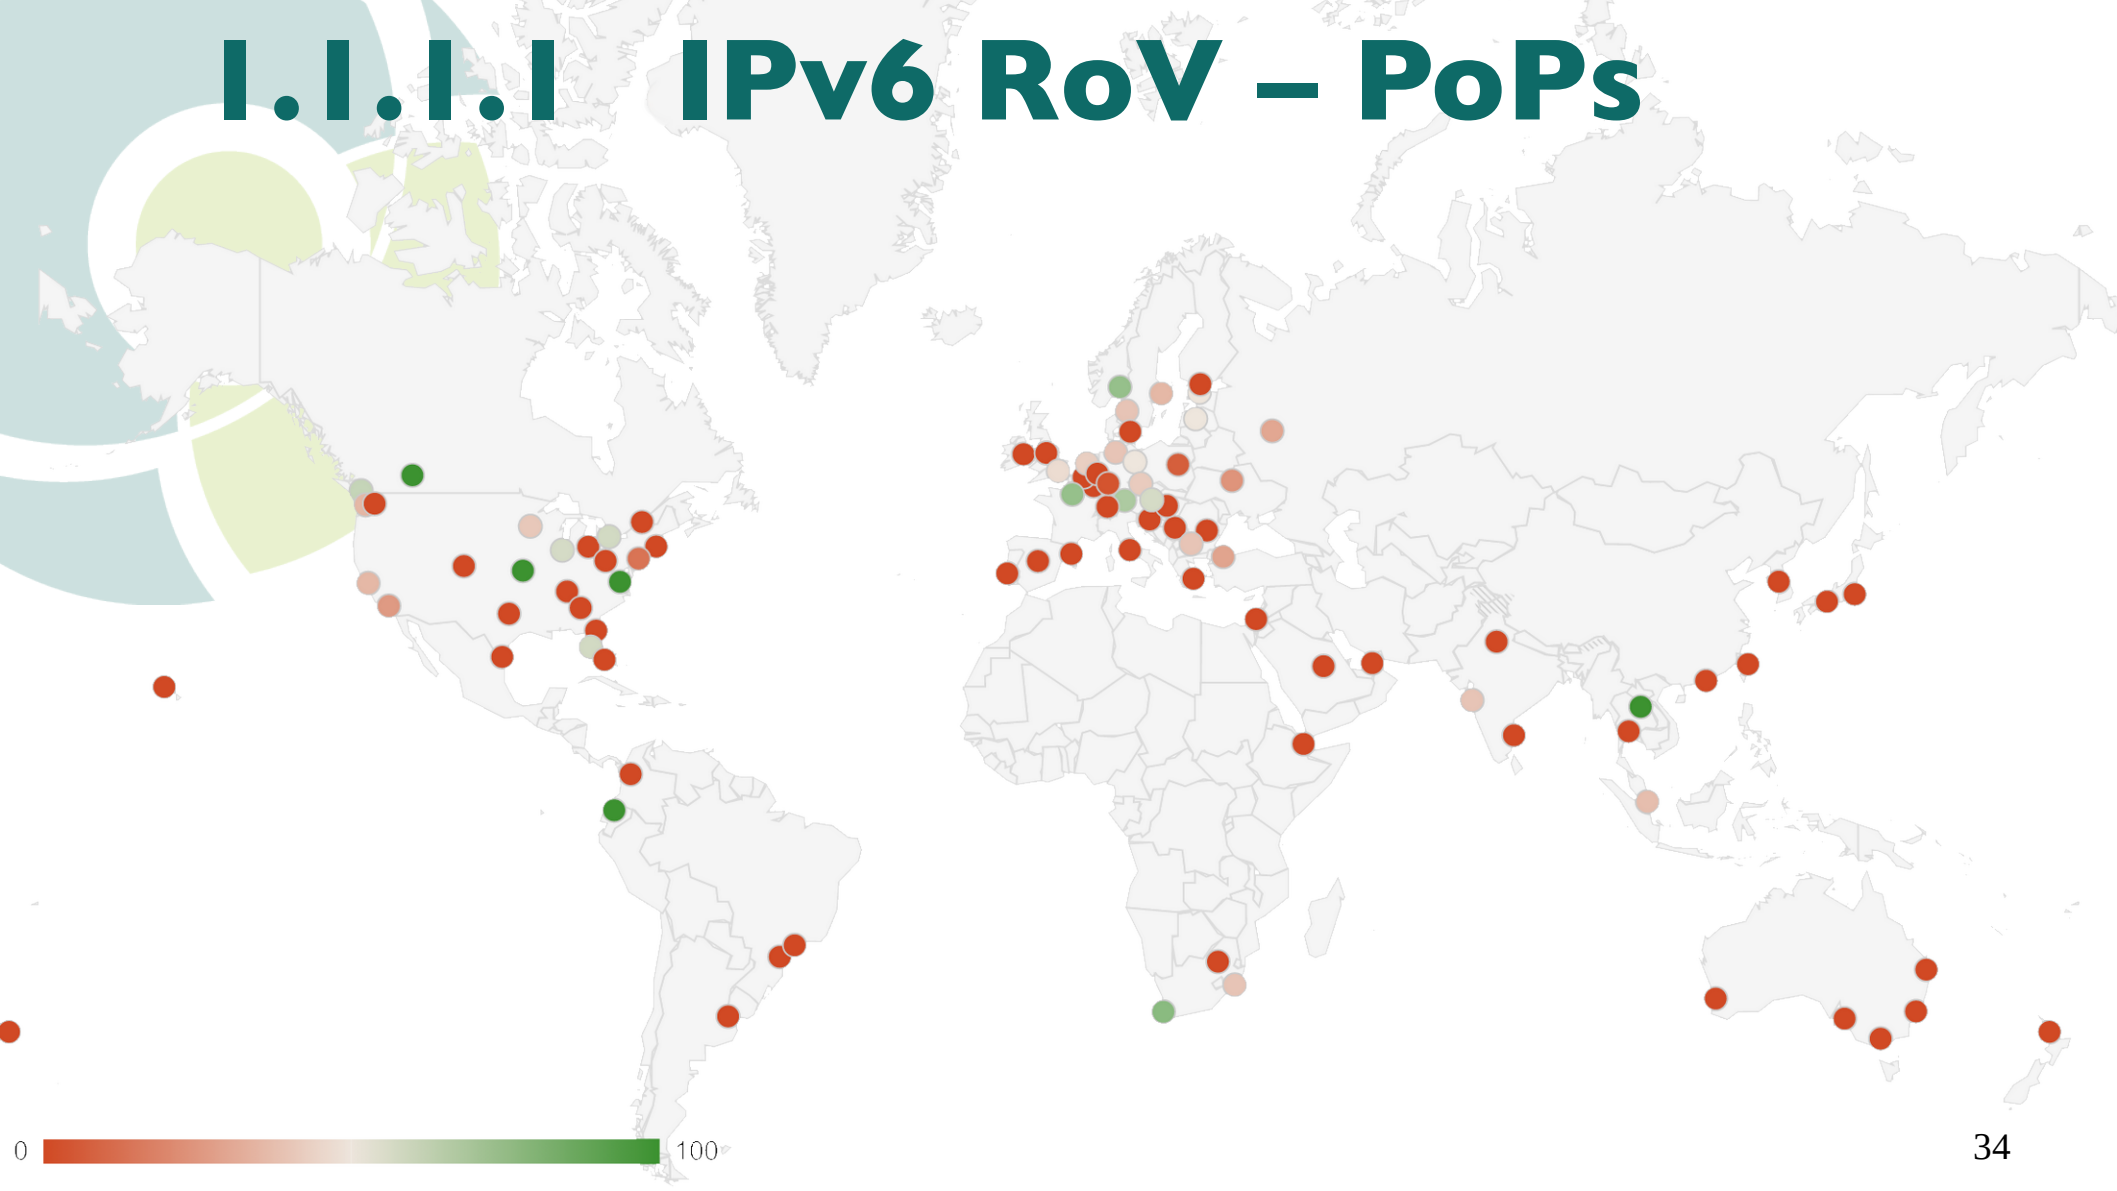

1.1.1.1 IPv6 RoV – Countries

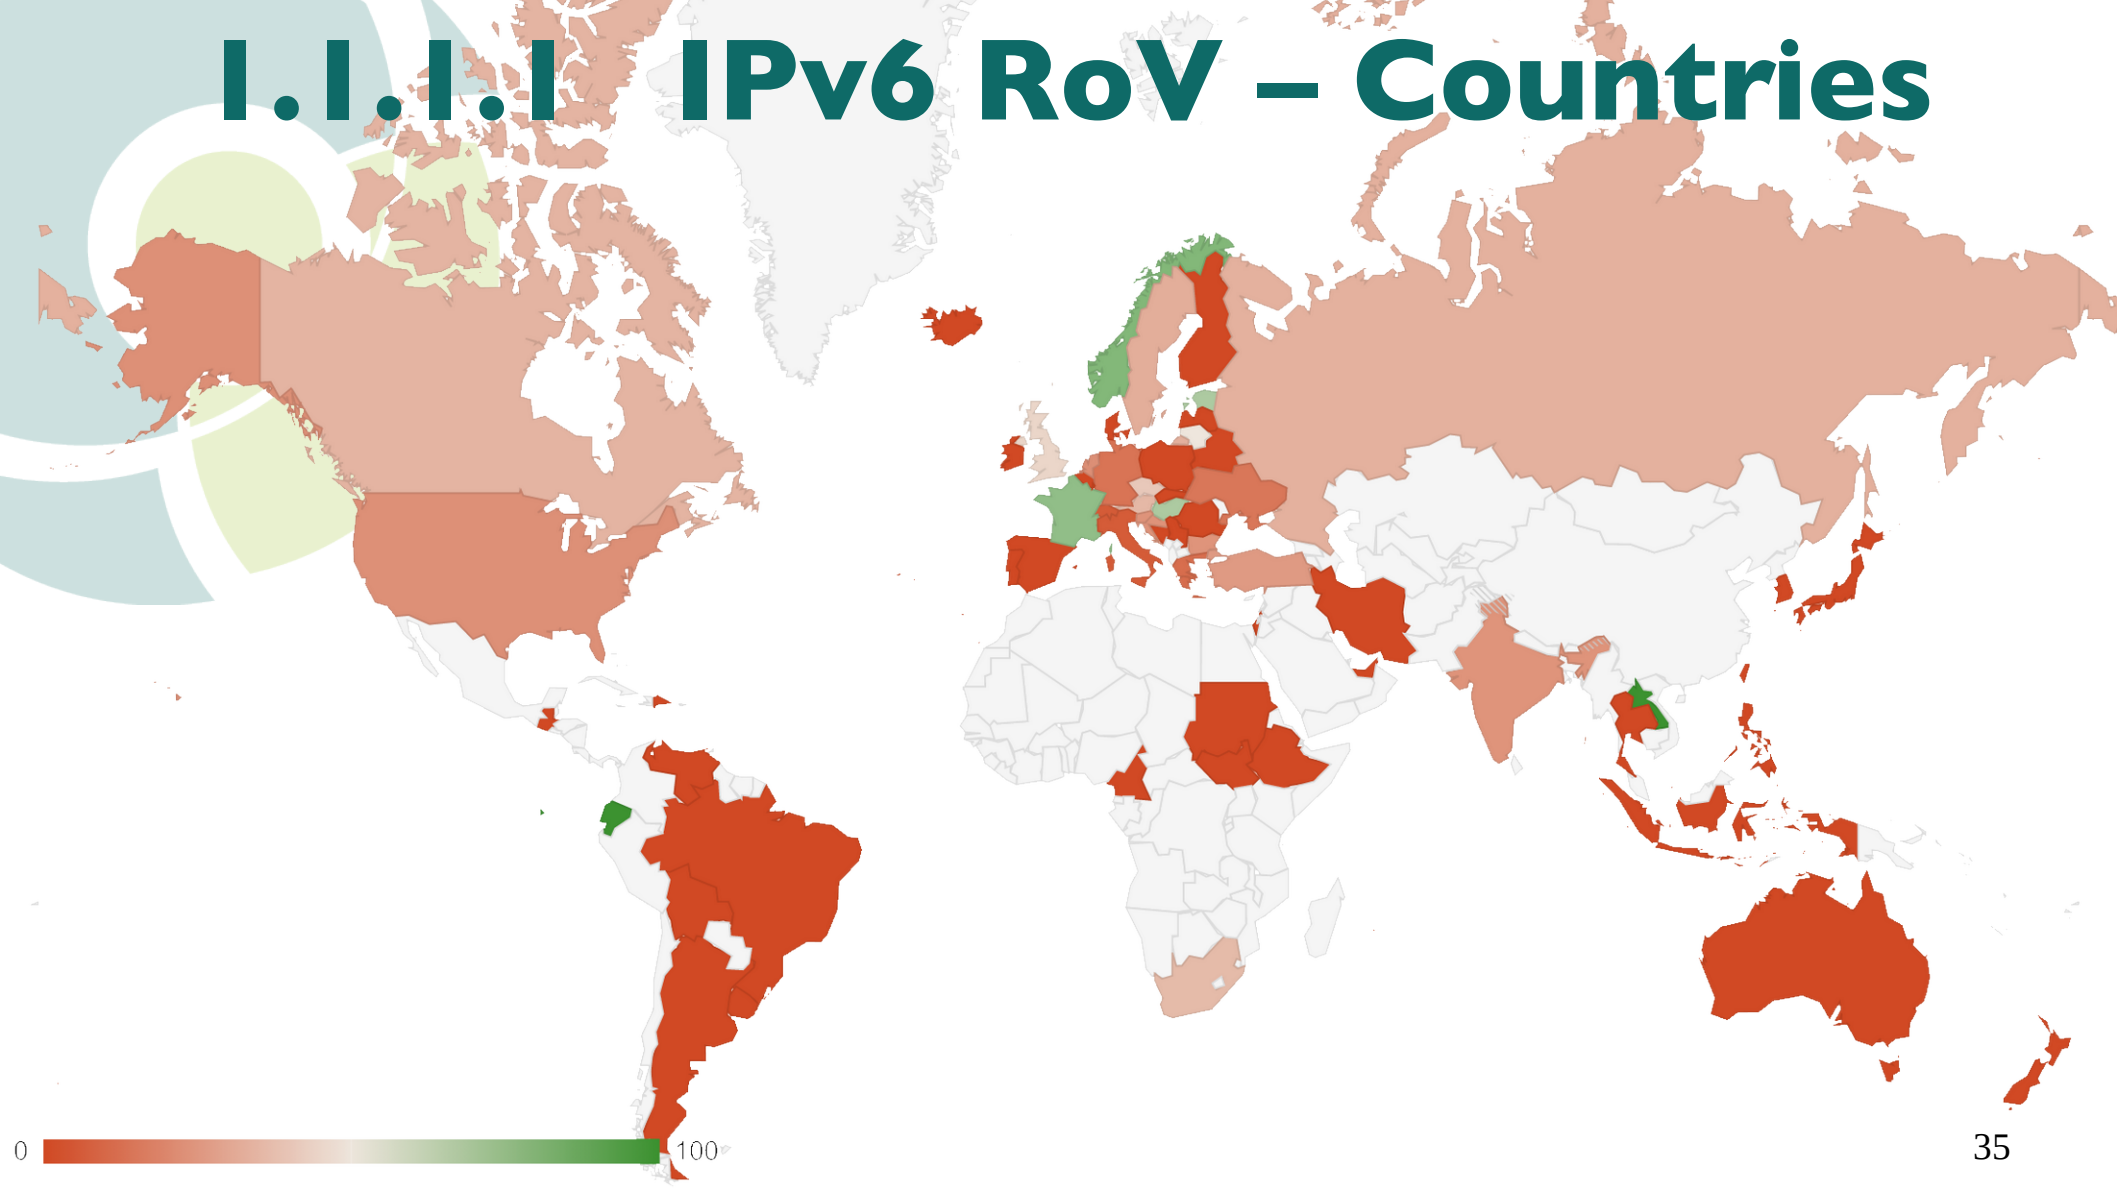

Future improvements

- We looked at authoritatives only
  - measurement network with **more vantage points**!
- Beacons all over the world

- dnsthought results for (probe, resolver, IP @ auth)
- dnsthought measurements for *not* answering auth to inventory IP @ auth for (probe, resolver)

Questions?

- Research performed by:
  - Erik Dekker <Erik.Dekker@os3.nl>
  - Marius Brouwer <mbrouwer@os3.nl>
- From
  - University of Amsterdam
- At
  - NLnetLabs
- On
  - January 2020
- Report:
  - https://delaat.net/rp/2019-2020/p04/report.pdf
- DNSThought:
  - https://dnsthought.nlnetlabs.nl/