


Measuring Route Origin Validation

Setting up a useful
RPKI Beacon

Willem Toorop

20 March 2022

IEPG at IETF113 Vienna



Picture from the Maelstrom Ascendant [book](#) and [album](#).

Copyright by: **Duncan Smith**

Why this presentation

- I had to setup an RPKI Beacon, but I'm not an RPKI expert
- I need your opinion – does what I did make sense
- I need your brains – how to best measure ROV?
- I want the beacon to be usable (available) for you too!
- We could also use some additional IPv4 resources 😊

Genesis

- I do DNS measurements (with RIPE Atlas mostly)
- Job Snijders offered to use his RPKI Beacon during IMC2019 in Amsterdam 🙏💖
- Started measuring the uptake Of Route Origin Validation since Januari 2020
- Job's beacon EOL in Oct. 2020



RPKI TEST a RIPE Labs experiment



```

testing valid ROA...[passed]
testing invalid ROA (5sec)...[passed]
AS2121 drops RPKI invalid BGP routes from prefix 2001:67c:64::/48 as witnessed
by your public IP 2001:67c:64:42:84ae:5b00:a129:e85c
  
```

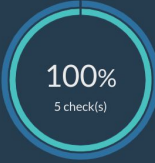
Check My DNS - Chromium

Check My DNS cmdns.dev.dns-oarc.net

Check My DNS

A

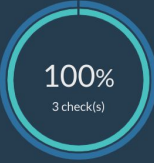
Transport



100%

5 check(s)

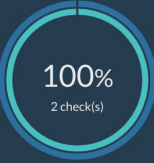
Basic DNS



100%

3 check(s)




DNS Features



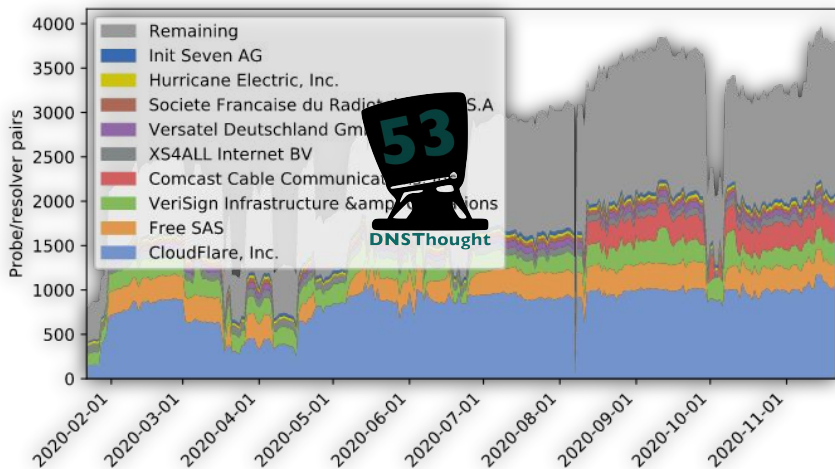
100%

2 check(s)

Achievements

© 2021 OARC, Inc. Feedback @ GitHub / @dnsoarc / email



The Current State of DNS Resolvers and RPKI Protection

Marius Brouwer
University of Amsterdam
marius.brouwer@os3.nl

Erik Dekker
University of Amsterdam
erik.dekker@os3.nl

ABSTRACT

The goal of this research was to gain insight into the Resource Public Key Infrastructure (RPKI) protection state of DNS resolvers. RIPE Atlas Probes were used to send DNS queries to an authoritative DNS server. This server contained Resource Records in both an RPKI valid and invalid prefix. The RIPE Atlas probes were instructed to send their queries


security, it is not broadly adopted [7, 8]. For this reason, this paper will focus on RPKI.

Due to the distributed nature of BGP and RPKI, the majority of network operators should sign their network prefixes and implement RPKI filtering to minimize prefix hijacks and route leaks [9]. A study conducted in 2019 claims that between 9.98% and 11.28% of the BGP announcements are verifiable using RPKI [10].

Re-evaluate setup - before

- Invalid IPv4 /24 & IPv6 /48 + /24 & /48 valids for reference
 - ✓ If endpoint validates → invalid = unreachable
 - ✗ If any hop in between validates → invalid = unreachable
 - ✗ Validating hop may be in return path

Re-evaluate setup - before



Re-evaluate setup - before

- Invalid IPv4 /24 & IPv6 /48 + /24 & /48 valids for reference
 - ✓ If endpoint validates → invalid = unreachable
 - ✗ If any hop in between validates → invalid = unreachable
 - ✗ Validating hop may be in return path
 - ✗ Is this a realistic route hijack?
 - ✗ Unreachable detection based on timeout

Re-evaluate setup - new setup

- Valid /23 (IPv4) and /47 (IPv6) and Invalid /24 and /48 more specific announcements from elsewhere
 - ✓ More realistic route hijack?
 - ✓ Don't have to wait for timeouts!

Re-e

- Vali
mor

The screenshot shows a web browser window titled "Krill - RPKI - Chromium" with the URL "prod-ca.krill.cloud/index.html#/cas/nlnetlabs". The page header features the Krill logo, a language dropdown set to "English", and user profile icons. The main content area is titled "Certificate Authority nlnetlabs" and has three tabs: "ROAs" (selected), "Parents", and "Repository". A search bar contains the IP address "185.49.142.0" and a "Download CSV" button is visible. Below the search bar is a table with columns for "ASN", "Prefix", and "State".

ASN	Prefix	State	
> 0	185.49.142.0/24-24	REDUNDANT	
> 14618	185.49.142.0/23-23	DISALLOWING 1	
> 16509	185.49.142.0/23-23	SEEN 1 1	
211321	185.49.142.0/24	INVALID ASN	


48

Re-evaluate setup - new setup

- Valid /23 (IPv4) and /47 (IPv6) and Invalid /24 and /48 more specific announcements from elsewhere
 - ✓ More realistic route hijack?
 - ✓ Don't have to wait for timeouts!
 - ✗ Still can't determine which hop is validating
 - ✗ **Even when your network is validating, you can still reach the invalid!**

Re-evalu

- Valid /23 (more spec
- ✓ More r
- ✓ Don't h
- x Still ca
- x Even v
- you ca



- and /48


e

g

Picture from the Maelstrom Ascendant [book](#) and [album](#).

Copyright by: **Duncan Smith**

Re-evaluate setup - new setup



Apps

- CMDNS:

Moved 👍

- Resolver test:

Very fast! 👍

- RPKI NCC

Web test:

Ont it's way 👍

```
willem@makaak: ~  
willem@makaak:~$ dig @1.1.1.1 rpkitest4.nlnetlabs.nl TXT  
  
; <<>> DiG 9.16.15-Ubuntu <<>> @1.1.1.1 rpkitest4.nlnetlabs.nl TXT  
; (1 server found)  
;; global options: +cmd  
;; Got answer:  
;; ->HEADER<<- opcode: QUERY, status: NOERROR, id: 36209  
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1  
  
;; OPT PSEUDOSECTION:  
; EDNS: version: 0, flags:; udp: 1232  
;; QUESTION SECTION:  
;rpkitest4.nlnetlabs.nl.                IN      TXT  
  
;; ANSWER SECTION:  
rpkitest4.nlnetlabs.nl. 1          IN      TXT      "HOORAY - Your resolver is prote  
cted by Route Origin Validation :)"  
  
;; Query time: 32 msec  
;; SERVER: 1.1.1.1#53(1.1.1.1)  
;; WHEN: do mrt 17 15:41:43 CET 2022  
;; MSG SIZE rcvd: 130  
  
willem@makaak:~$
```

For you!



@ RIPE Atlas

The screenshot shows a web browser window with the URL `atlas.ripe.net/probes/1003688/#tab-network`. The page title is "Probe on RPKI Invalid resources". The main content is divided into two columns for IPv4 and IPv6 configurations. The IPv4 section shows a current configuration with an Internet Address of `185.49.142.11`, ASN `211321 (NLNETLABS - Stichting NLnet Labs)`, Local Address `185.49.142.11`, Gateway `185.49.142.1`, Netmask `255.255.255.0`, and DNS_Resolvers `185.49.142.1`. The IPv6 section shows a current configuration with Addresses `2a04:b904::11/64` and `2a04:b907::11/64`, ASN `211321 (NLNETLABS - Stichting NLnet Labs)`, Gateway(s) `2a04:b907::1` and `2a04:b904::1`, and DNS_Resolvers `2a04:b907::1`. Below these is a "Probe Address Discovery" table. On the right side, there is a sidebar with the probe ID `#1003688`, a "1 week" timer, and a map of the location in Hembrug, Amsterdam, with a green location pin.

Probe on RPKI Invalid resources

General Network Built-ins UDMs Status (beta)

IPv4

Current Configuration

Internet Address	185.49.142.11
ASN	211321 (NLNETLABS - Stichting NLnet Labs)
Local Address	185.49.142.11
Gateway	185.49.142.1
Netmask	255.255.255.0
DNS_Resolvers	185.49.142.1

IPv6

Current Configuration

Addresses	2a04:b904::11/64 2a04:b907::11/64
ASN	211321 (NLNETLABS - Stichting NLnet Labs)
Gateway(s)	2a04:b907::1 2a04:b904::1
DNS_Resolvers	2a04:b907::1

Probe Address Discovery

	IPv4	IPv6
Connection Address	-	✓ 2a04:b904::11
IP Echo Service	-	-
The Local IP	✓ 185.49.142.11	2a04:b907::11

Update Location

For you!



@ RIPE Atlas

@ NLNOG Ring

```
nlnetlabs@nlnetlabs01: ~
nlnetlabs@nlnetlabs01: ~ 89x33
NLNOG
RING Project

Welcome on nlnetlabs01.ring.nlnog.net, an NLNOG RING Node!
System operated by NLnet Labs - tech-admin@nlnetlabs.nl
Location: Netherlands - AS211321

Munin:

http://munin.infra.ring.nlnog.net/munin/ring.nlnog.net/nlnetlabs01.ring.nlnog.net/

For more information, please visit https://ring.nlnog.net/

0 updates can be applied immediately.

Last login: Sun Mar 20 07:34:20 2022 from 2a04:b900::1:0:0:10

***[ RPKI Beacon ]*****
**
** This NLNOG Ring Node contains an extra alternative network **
** namespace which has IP resources which are RPKI Invalid on **
** purpose. To enter this alternative network namespace, use: **
**
**   enter-invalid-netns [program [arguments]] **
**
*****
nlnetlabs@nlnetlabs01:~$
```

For you!



@ RIPE Atlas

@ NLNOG Ring

```
nlnetlabs@nlnetlabs01: ~
nlnetlabs@nlnetlabs01: ~ 89x33
** This NLNOG Ring Node contains an extra alternative network namespace which has IP resources which are RPKI Invalid on purpose. To enter this alternative network namespace, use:
**
**      enter-invalid-netns [program [arguments]]
**
*****
nlnetlabs@nlnetlabs01:~$ ping www.ietf.org
PING www.ietf.org(2606:4700::6810:2c63 (2606:4700::6810:2c63)) 56 data bytes
64 bytes from 2606:4700::6810:2c63 (2606:4700::6810:2c63): icmp_seq=1 ttl=60 time=2.96 ms
64 bytes from 2606:4700::6810:2c63 (2606:4700::6810:2c63): icmp_seq=2 ttl=60 time=16.7 ms
64 bytes from 2606:4700::6810:2c63 (2606:4700::6810:2c63): icmp_seq=3 ttl=60 time=1.69 ms
^C
--- www.ietf.org ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 1.695/7.146/16.775/6.828 ms
nlnetlabs@nlnetlabs01:~$ enter-invalid-netns
nlnetlabs@nlnetlabs01:~$ ping www.ietf.org
PING www.ietf.org(2606:4700::6810:2d63 (2606:4700::6810:2d63)) 56 data bytes
^C
--- www.ietf.org ping statistics ---
4 packets transmitted, 0 received, 100% packet loss, time 3060ms

nlnetlabs@nlnetlabs01:~$ ping www.nlnetlabs.nl
PING www.nlnetlabs.nl(dicht.nlnetlabs.nl (2a04:b900::1:0:0:10)) 56 data bytes
64 bytes from dicht.nlnetlabs.nl (2a04:b900::1:0:0:10): icmp_seq=1 ttl=58 time=1.61 ms
64 bytes from dicht.nlnetlabs.nl (2a04:b900::1:0:0:10): icmp_seq=2 ttl=58 time=1.58 ms
64 bytes from dicht.nlnetlabs.nl (2a04:b900::1:0:0:10): icmp_seq=3 ttl=58 time=1.59 ms
^C
--- www.nlnetlabs.nl ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 1.587/1.601/1.619/0.013 ms
nlnetlabs@nlnetlabs01:~$
```


For you!



@ RIPE Atlas


@ NLNOG Ring

Measure services
Should they?

@ Anything else?

```
nlnetlabs@nlnetlabs01: ~  
nlnetlabs@nlnetlabs01:~$ enter-invalid-netns  
nlnetlabs@nlnetlabs01:~$ dig @a0.org.afiliias-nst.info. ietf.org  
  
; <<>> DiG 9.11.3-1ubuntu1.17-Ubuntu <<>> @a0.org.afiliias-nst.info. ietf.org  
; (2 servers found)  
;; global options: +cmd  
;; Got answer:  
;; ->HEADER<- opcode: QUERY, status: NOERROR, id: 2943  
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 6, ADDITIONAL: 1  
;; WARNING: recursion requested but not available  
  
;; OPT PSEUDOSECTION:  
; EDNS: version: 0, flags:; udp: 1232  
;; QUESTION SECTION:  
;ietf.org.                IN      A  
  
;; AUTHORITY SECTION:  
ietf.org.                86400  IN      NS      ns1.sea1.afiliias-nst.info.  
ietf.org.                86400  IN      NS      ns0.amsl.com.  
ietf.org.                86400  IN      NS      ns1.hkg1.afiliias-nst.info.  
ietf.org.                86400  IN      NS      ns1.mia1.afiliias-nst.info.  
ietf.org.                86400  IN      NS      ns1.yyz1.afiliias-nst.info.  
ietf.org.                86400  IN      NS      ns1.ams1.afiliias-nst.info.  
  
;; Query time: 1 msec  
;; SERVER: 2001:500:e::1#53(2001:500:e::1)  
;; WHEN: Sun Mar 20 07:54:46 UTC 2022  
;; MSG SIZE rcvd: 194  
  
nlnetlabs@nlnetlabs01:~$
```

Why - questions & feedback

- I need your opinion – **does what I did make sense?**
- I need your brains – **how to best measure ROV?**
- I want the beacon to be usable (available) for you too!
– **what tool do you want/need?**
- We could also use some additional IPv4 resources 
– **Collaborate?**