# 

Long Term Analysis of Root Server System Performance using RIPE Atlas Data Duane Wessels

31<sup>st</sup> DNS-OARC Workshop, Austin TX

November 1, 2019

# Outline

- RIPE Atlas & Its Data
- Root Server Operator (RSO) Availability
- RSO Latency
- RSO Interception
- Root Server System (RSS) Availability
- RSS Latency
- Coverage



### **RIPE** Atlas

- <u>https://atlas.ripe.net/</u>
- 10,000+ probes connected
- Operating since 2011
- Coverage biased toward RIPE region (Europe)



About 550 Atlas anchors





# Built-In "hostname.bind" Measurements

- Measurement ID numbers 10301-10316, 11301-11316
- To root server IP addresses
- IPv4 and IPv6
- Every 4 minutes
- Provides RTT to all root servers
- Also provides data on DNS interception



Built-in Random Domain Name (TLD) Measurement

- Measurement 30001
- Since Feb 2017
- Every 10 minutes
- Query sent to probe's configured resolver
- Measure RTT to "root server system"



### **Data Collection**



# **Data Collection**

- Want to have one measurement per probe, per root server address, per day.
- Measurement times chosen randomly each day
- Some 323,000,000 individual measurements from 2012 to 2019.



### Example measurement

```
{
        "af": 4,
        "dst_addr": "198.41.0.4",
        "dst port": "53",
        "from": "193.0.0.78",
        "fw": 4945,
        "lts": 8,
        "msm id": 10309,
        "msm name": "Tdig",
        "prb id": 10004,
        "proto": "UDP",
        "result": {
            "ANCOUNT": 1,
            "ARCOUNT": 0,
            "ID": 15968,
            "NSCOUNT": 0,
            "ODCOUNT": 1,
            "abuf":
"PmCEAAABAAEAAAAACGhvc3RuYW11BGJpbmQAAB
AAA8AMABAAAwAAAAAACglubm4xLWxvbjg=",
```

```
"answers": [
        {
            "NAME": "hostname.bind",
            "RDATA": [
              "nnn1-lon8"
            ],
            "TYPE": "TXT"
        }
    1,
    "rt": 9.402,
    "size": 53
},
"src addr": "193.0.3.34",
"stored timestamp": 1535831672,
"timestamp": 1535831575,
"type": "dns"
```

},

### Number of Atlas Probes Over Time



### Number of IPv4 Measurements Per Day



### Number of IPv6 Measurements Per Day







# **Calculating Availability**

· Most measurements provide a successful result.

### • But some result in errors.

- timeout 5000
- socket connect failed Network is unreachable
- senderror AF Network is unreachable, NOT AF\_INET
- senderror AF Network is unreachable, AF\_INET
- senderror Network is unreachable
- socket connect failed Permission denied
- socket connect failed Network unreachable
- timeout 1000
- getaddrinfo port domain, AF 10 Servname not supported for ai\_socktype
- evdns\_getaddrinfo U nodename nor servname provided, or not known
- getaddrinfo port domain, AF 2 Servname not supported for ai\_socktype

# $Availability = \frac{Count \ of \ success}{Count \ of \ success \ and \ timeout}$



### Availability, One RSO, All Probes



Atlas built-in measurements

### Availability, All RSOs, All Probes























### In no particular order

### IPv6 has lower availability?

- Atlas reports some "facts" about probes, such as
  - has working IPv4
  - has working IPv6
- What if we only include "working" probes?
- What about anchors?



### Availability, All RSOs, Working Probes





















### In no particular order

18

### Availability, All RSOs, Anchor Probes



100 80 60 40 20 2012 2013 2014 2015 2016 2017 2018 2019





















In no particular order

### Availability, One RSO, Anchor Probes



Atlas built-in measurements





### Median Latency, Individual RSO



### Median Latency, All RSOs





2012 2013 2014 2015 2016 2017 2018 2019















2012 2013 2014 2015 2016 2017 2018 2019



Source: derived from RIPE Atlas built-in measurements

23

### powered by VERISIGN

In no particular order

# Trend

- Most RSOs have a slight but noticeable downward trend over time.
- Is median latency getting better as more RSO instances are deployed?
- Or is it due to more RIPE Atlas probes over time?
- Let's look at latency from "the best" probes.
  - those which most consistently provide data.
- Also from Atlas anchors.



### Median Latency, "Best Probes"



<sup>200</sup> 150 50 0 2012 2013 2014 2015 2016 2017 2018 2019

















### In no particular order

Source: derived from RIPE Atlas built-in measurements

25

### powered by VERISIGN

### Median Latency, Anchor Probes



![](_page_24_Figure_2.jpeg)

2012 2013 2014 2015 2016 2017 2018 2019

![](_page_24_Figure_4.jpeg)

![](_page_24_Figure_5.jpeg)

![](_page_24_Figure_6.jpeg)

![](_page_24_Figure_7.jpeg)

2012 2013 2014 2015 2016 2017 2018 2019

![](_page_24_Figure_9.jpeg)

![](_page_24_Figure_10.jpeg)

![](_page_24_Figure_11.jpeg)

![](_page_24_Figure_12.jpeg)

![](_page_24_Figure_13.jpeg)

Source: derived from RIPE Atlas built-in measurements

Verisign Public

In no particular order

26

### Median Latency, All / Best / Anchor Probes

![](_page_25_Figure_1.jpeg)

![](_page_26_Figure_0.jpeg)

![](_page_26_Picture_1.jpeg)

## **Measuring Interception**

- · Atlas query is for "hostname.bind".
- Mismatch against known patterns implies interception, of some kind.

aa.ec-uio.l.root aa.pe-cll.l.root aa.tn-nbe.l.root dns2 aa.lb-bfz.l.root aa.us-hnl.l.root aa.cr-sjo.l.root aa.ru-yek.l.root aa.gu-hgt.l.root aa.py-asu.l.root reza-ProLiant-DL380-Gen9 aa.us-dsm.l.root aa.br-crn.l.root XDR on cdns11.enbqk aa.ma-rba.l.root ab.lk-rml.l.root

### Interception, Individual RSO

![](_page_28_Figure_1.jpeg)

### Interception, All RSOs

1

0.8

0.6

0.4

0.2

0

1

0.8

0.6

0.4

0.2

0

1

0.8

0.6

0.4

0.2

0

2012 2013 2014 2015 2016 2017 2018 2019

2012 2013 2014 2015 2016 2017 2018 2019

2012 2013 2014 2015 2016 2017 2018 2019

![](_page_29_Figure_1.jpeg)

In no particular order

![](_page_29_Figure_3.jpeg)

31

0.8

0.6

0.4

0.2

1

0.8

0.6

0.4

0.2

0

0.8

0.6

0.4

0.2

2012 2013 2014 2015 2016 2017 2018 2019

2012 2013 2014 2015 2016 2017 2018 2019

2012 2013 2014 2015 2016 2017 2018 2019

powered by VERISIGN

0.8

0.6

0.4

0.2

0

1

0.8

0.6

0.4

0.2

0

1

0.8

0.6

0.4

0.2

0

2012 2013 2014 2015 2016 2017 2018 2019

2012 2013 2014 2015 2016 2017 2018 2019

2012 2013 2014 2015 2016 2017 2018 2019

Verisign Public

### Root Server System Availability

Using Built-in Random Domain Name (TLD) Measurement

![](_page_30_Picture_2.jpeg)

![](_page_30_Picture_3.jpeg)

### **Root Server System Availability**

![](_page_31_Figure_1.jpeg)

33

### Root Server System Latency

### Using Built-in Random Domain Name (TLD) Measurement

![](_page_32_Picture_2.jpeg)

![](_page_32_Picture_3.jpeg)

### **Root Server System Median Latency**

![](_page_33_Figure_1.jpeg)

35

### Root Server System Coverage

How does the number of probes relate to coverage and measurement accuracy?

![](_page_34_Picture_2.jpeg)

![](_page_34_Picture_3.jpeg)

### Number of Probes and RSS Instances

![](_page_35_Figure_1.jpeg)

### Number of Probes and Latency Accuracy

![](_page_36_Figure_1.jpeg)

38

### Number of Probes and Latency Accuracy

![](_page_37_Figure_1.jpeg)

In no particular order

# In Summary...

![](_page_38_Picture_1.jpeg)

# **Summary of Typical Results**

- RSO availability: ~98% v4, ~95% v6
- RSS availability: ~98%
- RSO median latency: 25-50 milliseconds
- RSS median latency: 25-30 milliseconds
- RSO interception: 0.8%

41

![](_page_39_Picture_7.jpeg)

# Questions?

![](_page_40_Picture_1.jpeg)

![](_page_41_Picture_0.jpeg)

© 2016 VeriSign, Inc. All rights reserved. VERISIGN and other trademarks, service marks, and designs are registered or unregistered trademarks of VeriSign, Inc. and its subsidiaries in the United States and in foreign countries. All other trademarks are property of their respective owners.