



# ILLUMINATING THE INTERNET: RESEARCHER'S ADVENTURES IN INTERNET MEASUREMENT


Raffaele Sommesse – University of Twente – RIPE 87 Student Event

# A couple of bits about me

- Postdoctoral Researcher at **University of Twente**, The (rainy) Netherlands
- From the (sunny) Naples, Italy
- Working on **DNS hygiene, resilience and security** (I will talk a lot about it).
- Doing Internet Measurements since my bachelor thesis :)
- Love to dismantle stuff to understand how they works.

# Internet

- Internet is one of the most complex network that the human species has built and is the critical element of our modern society.
- It is an interconnection of diverse networks and governed by RFCs and best practices.
- Understanding global-scale interactions can be puzzling.



# Internet Stability




Internet's stability often averts minor issues.



Hard failures can lead to catastrophic consequences.



Measuring helps us grasp its resilience dynamics.



# Why do Internet measurements?


The Sherlock Holmes of the Digital World

- Monitor and record internet history.
- Detect potential threats and prevent them.
- Identify RFC and best practices non-compliance and their consequences.
- Uncover outages, censorship, and failure events.

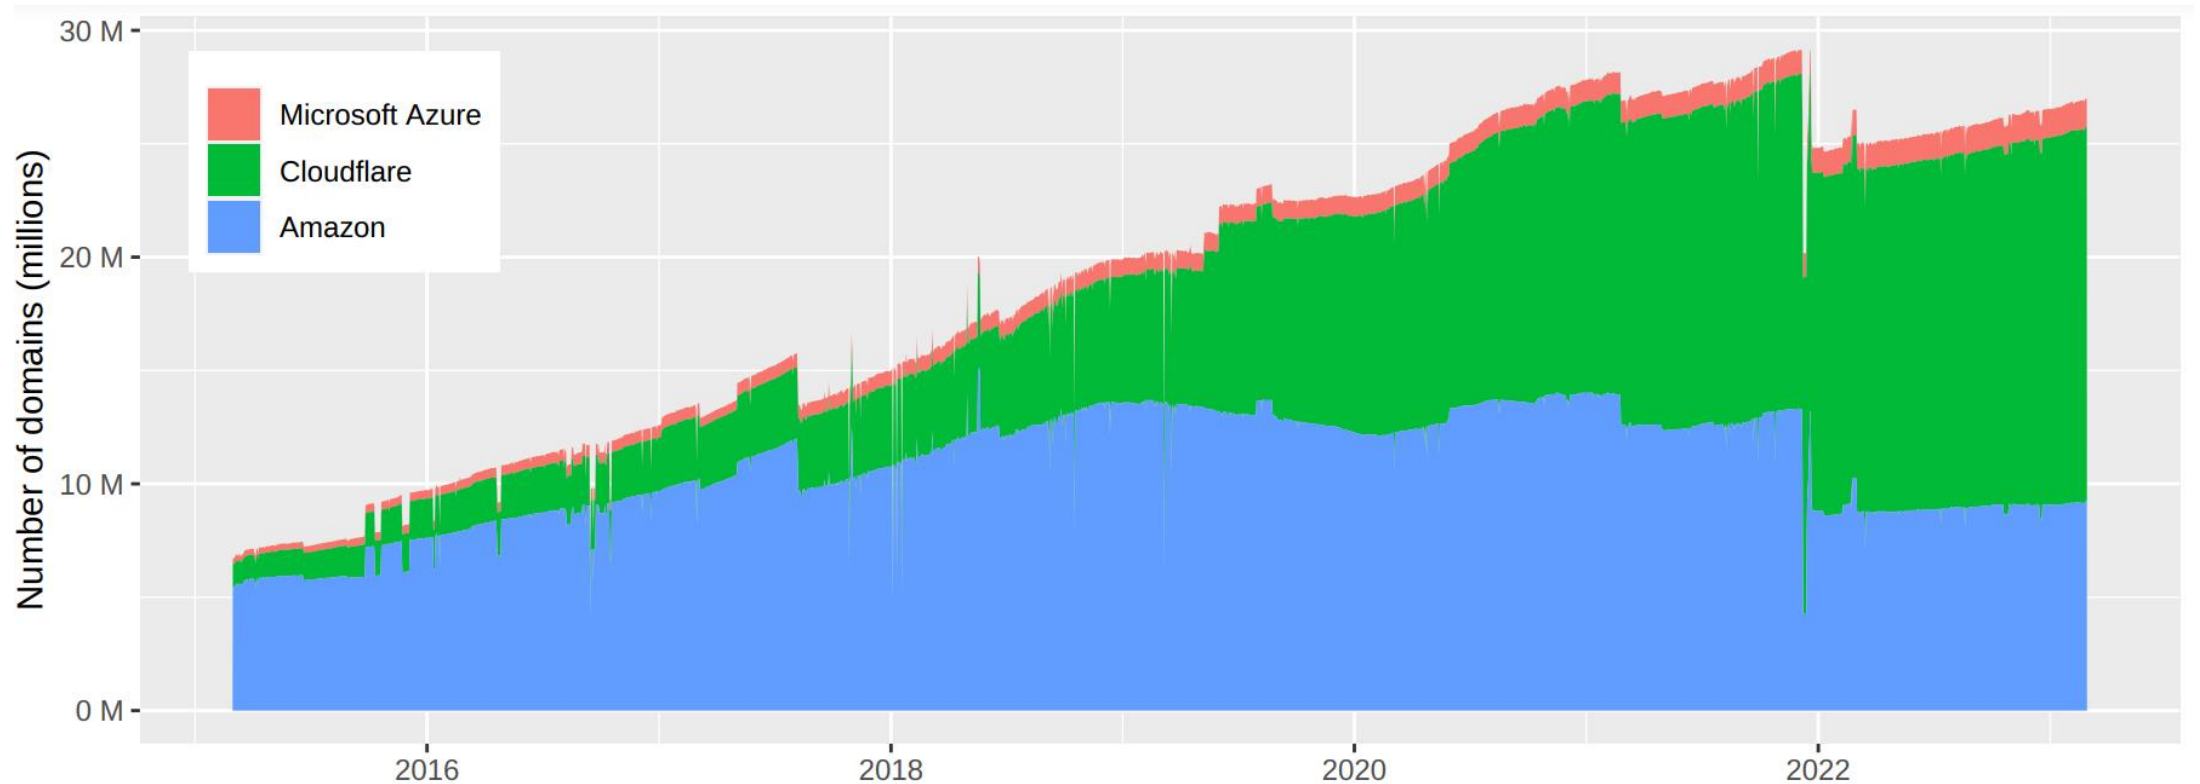


# Internet History (1)

- Observe how it evolved in different regions and with our society.
- Imagine history books for the internet!
- We can answer to questions like: what was the DNS server of the New York Time in 2012?

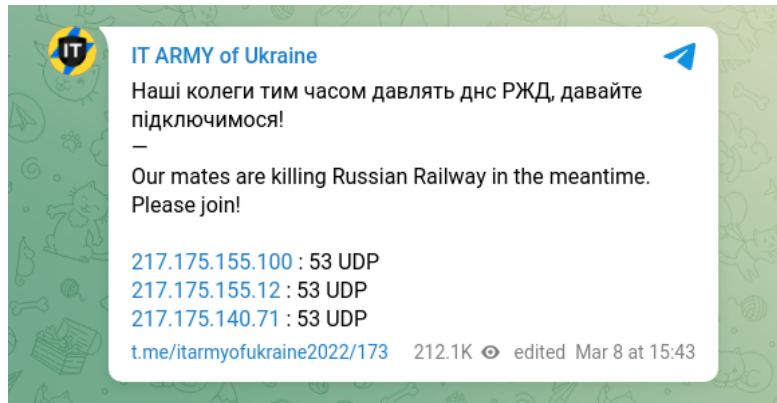


# Internet History (2)

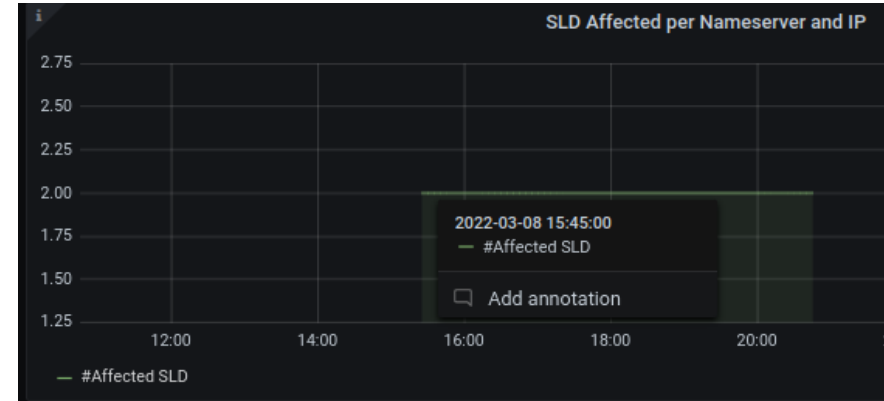


Cloud adoption evolution 2015-2023. Source: OpenINTEL


# Uncovering outages and attacks (1)



Telegram



Attack detected by UCSD Network






Reactive DNS Measurements



# Uncovering outages and attacks (2)

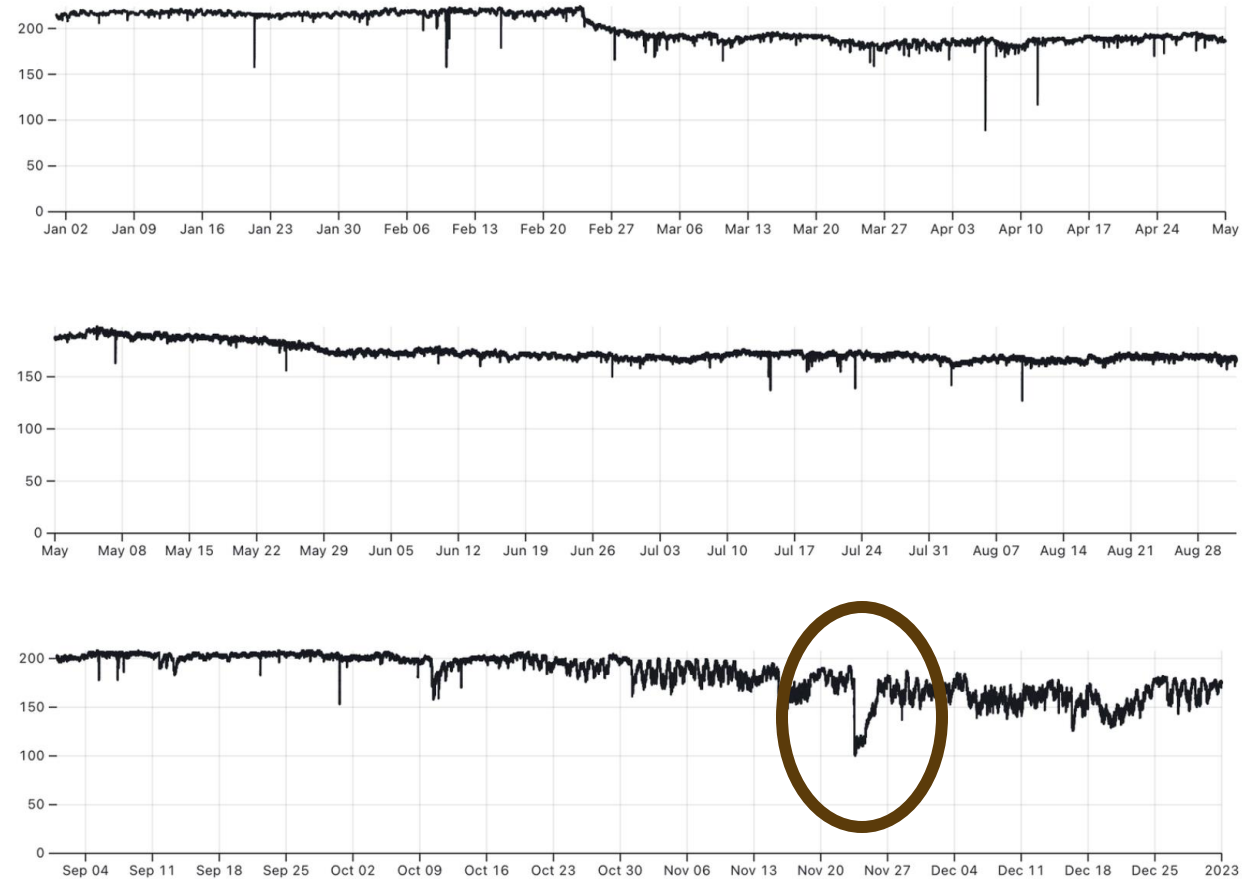
Ukrainian RIPE Atlas probe status: 2022-2023



Connected probes in 2022 in UA




# Uncovering outages and attacks (2)

Ukrainian RIPE Atlas probe status: 2022-2023



Connected probes in 2022 in UA

# Uncovering outages and attacks (2)



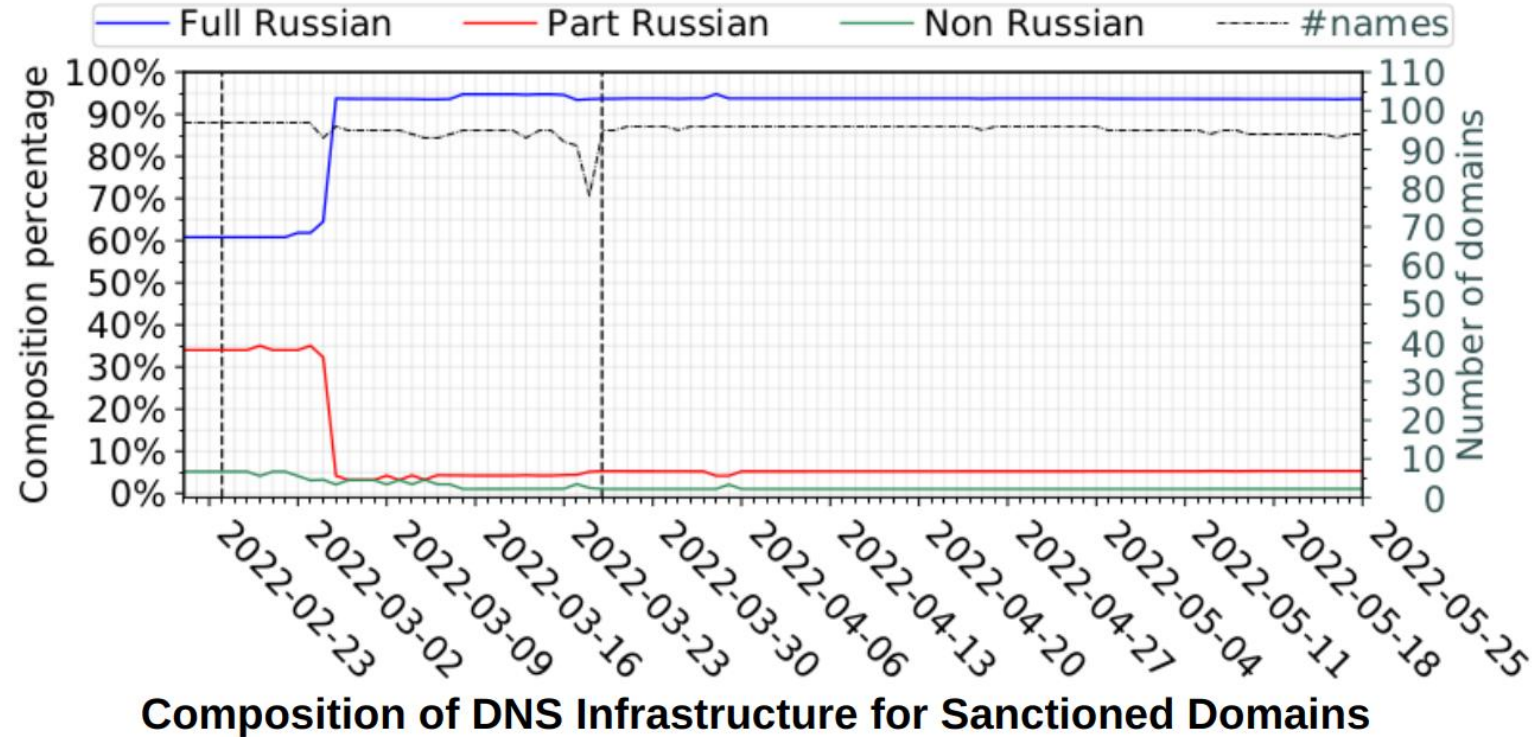
Connected probes in 2022 in UA

## Ukraine battles to restore power as millions face blackouts

*With temperatures falling below zero, authorities struggle to restore power and water supply as millions are cut off from the electricity grid.*

# Mapping sovereignty and sanctions

- Exodus of service providers from Russia.

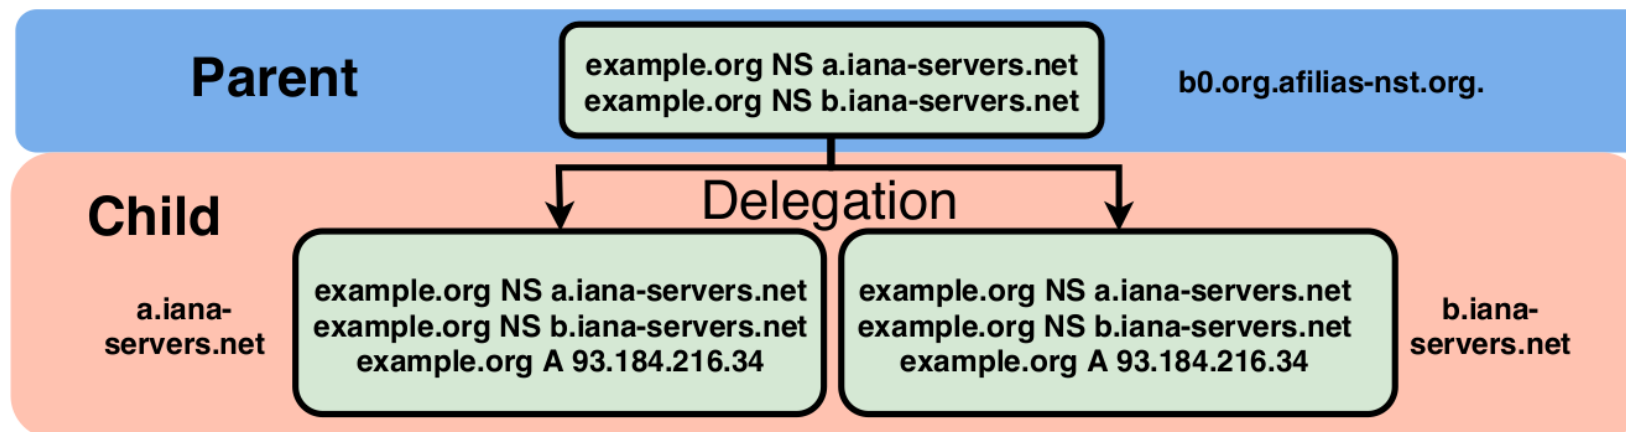


# — Are we all playing the same rules?

- Internet measurements can be used (and are widely used) to verify the compliance of software or operators to best practices and RFCs.
- With them we can identify spread and size of misconfiguration that may affect the stability and the functionality of the global Internet.
-

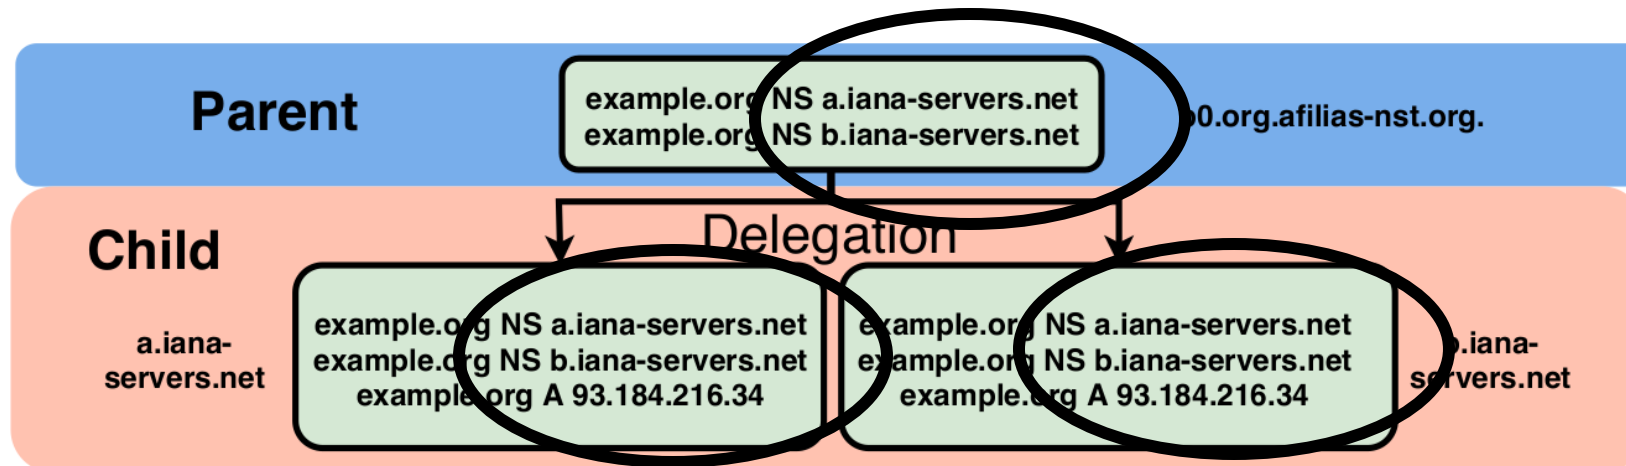
# An RFC non-compliance example

- The administrators of both (parent and children n.d.r.) zones should insure that the NS and glue RRs which mark both sides of the cut are consistent and remain so. (RFC 1034 –November 1987)




# An RFC non-compliance example

- The administrators of both (parent and children n.d.r.) zones should insure that the NS and glue RRs which mark both sides of the cut are consistent and remain so. (RFC 1034 –November 1987)



# Not always consistent!



.com: 11 million domains  
.net: 1 million domains  
.org: 0.7 million domains

Total: **12.7 million domains!**

Including popular domain like the **Indian TLD** and **att.com** (fixed after our notification)

NS inconsistency ( $P \neq C$ ) from 2017-04-01 until 2019-10-01

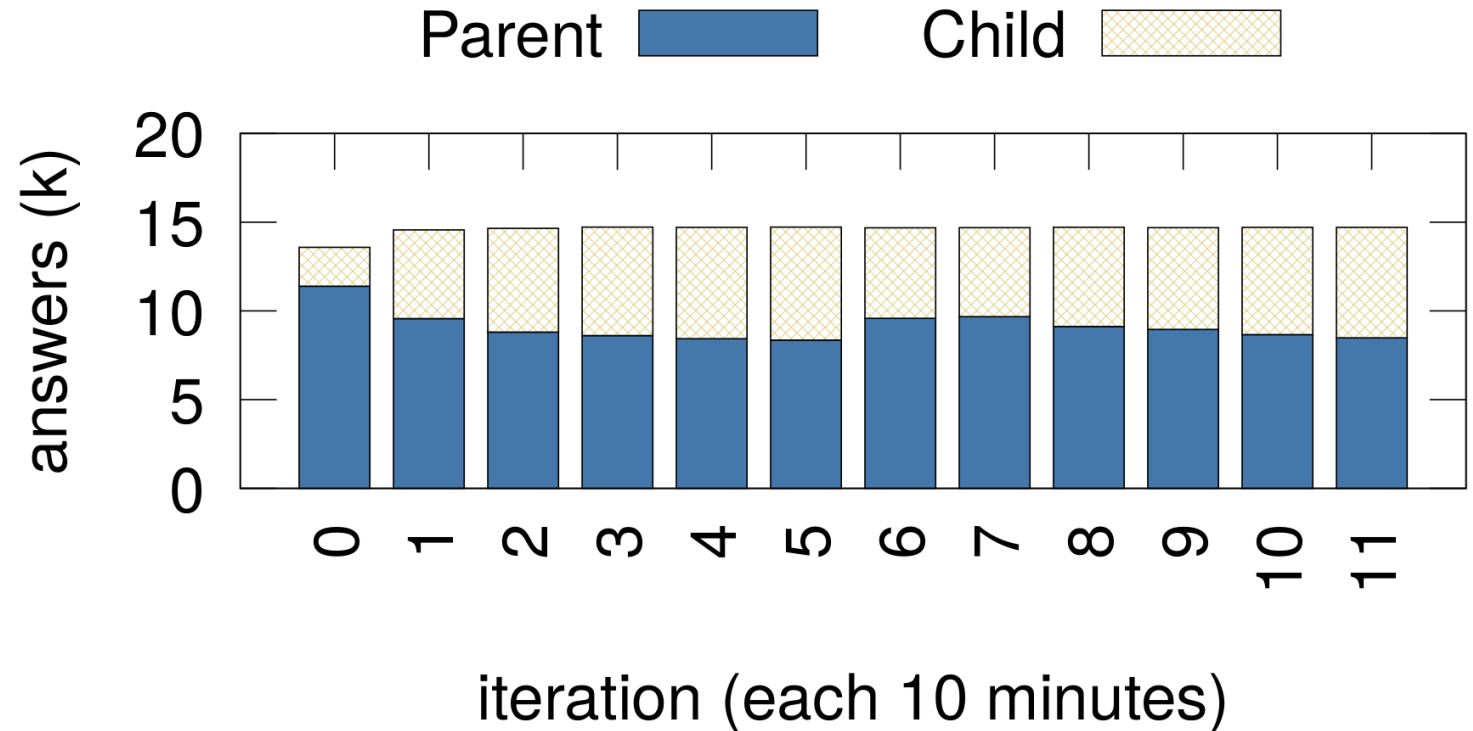


# Not even the resolver's behavior

RIPE Atlas

Measurements of  
an inconsistent  
domain

Clients may get  
different  
responses based  
on their resolver  
configuration



# Internet measurement toolkit

From your own PC:

- Standard: Ping (latency), Traceroute (network path), Dig (DNS), Nmap (Port scan), etc.
- Hardcore: Zmap, MassDNS, MassScan (be wise!)

Distributed measurements platform:

- Ripe ATLAS, Ark, Looking Glass, etc.

Publicly available datasets:

- Ripe RIS, Routeviews, Atlas Measurements, OpenINTEL, DZDB etc.

# Lots of data

For a bit of context, in OpenINTEL we scan 63% of the DNS SLDs space collecting:

- 252 million domains per day
- 4.3 billion data points per day
- 9.1 trillion data points since 2015
- More than 100 GB of compressed data every day
- More than 200 TB of data since 2015

Lots of challenge for data processing (and understanding!)

# A dumpster fire

- "Never attribute to malice that which is adequately explained by stupidity" - Hanlon's razor
- Analyzing data and discerning trends, threats, vulnerabilities and outages can be challenging due to the **heterogenous** variety of devices, people and **operators** playing on the global Internet.



# **A scientific approach**

- Correlation is not causation!
- Always trying to corroborate insights obtained from one dataset with data from other datasets!

# Conclusion

- Despite the challenges: Internet measurement is (a lot of) fun!
- Both the academic world and industry invest significant resources in Internet measurements studies.
- Thinking of a career in the field (?)

# Questions (?)

[r.sommese@utwente.nl](mailto:r.sommese@utwente.nl)

Hope to see you all in  
RIPE 87!

